

Fusion action systems

Matthew Gelvin

August 10, 2010

0 Introduction

Let p be a prime.

The theory of fusion systems is an abstraction of the p -local structure of finite groups. First codified by Puig in language that can be found in [Pui1], fusion systems have since become a subject of interest to both group theorists and block theorists. There are also applications to algebraic topology, originally in search of a solution to the Martino-Priddy Conjecture, which leads to the theory of p -local finite groups or classifying spaces of fusion systems, as introduced in [BLO2]. Much of the recent work on these subjects uses a different, though equivalent, formalism to describe fusion systems; a good source for this point of view is [Lin].

The work of this paper, which was mostly done toward completion of the author's doctoral thesis, began with an investigation of what it would mean for a fusion system to act on a finite set. An initial answer took the form of a fusion system \mathcal{F} on a p -group S , together with a finite set X with an S -action that “respected fusion data.” These \mathcal{F} -stable S -sets appear in [BLO2], where they are used to compute the cohomology ring of \mathcal{F} . From the topological point of view, a finite G -set gives rise to a covering space of BG ; the hope that a similar situation would realize for p -local finite groups.

Ultimately, it became apparent that simply imposing a condition on an \mathcal{F} -set did not yield enough structure to do, for example, unstable homotopy theory. To rigidify this flabbiness, we returned to “ambient case” of a finite group G acting on the set X with restricted action of $S \in \text{Syl}_p(G)$ to see what structure had been lost. By examining the simultaneous action of G on X and S (where the latter is given by conjugation, and is the same sort of “action” that underlies fusion theory in general), we were led to define the notion of *fusion action system*, and later *p -local finite group actions*.

0.1 Results

Explicitly, given a finite group G acting on the finite set X and $S \in \text{Syl}_p(G)$, the fusion action system relative to these data is the category \mathfrak{X}_G whose objects are the subgroups of S and whose morphisms are given by

$$\mathfrak{X}_G(P, Q) = \{(\varphi, \sigma) \mid \varphi \in \text{Inj}(P, Q), \sigma \in \Sigma_X, \text{ and } \exists g \in G \text{ s.t. } (\varphi, \sigma) = (c_g|_P, \ell_g)\}.$$

Here, Σ_X denotes the symmetric group on X , $c_g : P \rightarrow Q$ is the injective morphism given by conjugation by g , and ℓ_g is the permutation of X induced by g . The p -local finite group action arising from this situation is \mathfrak{X}_G together with a more complicated category that can be roughly thought of as restoring to the fusion action system the data of which elements of G induce a morphism in \mathfrak{X}_G . Much of this paper is devoted to properly abstractifying these concepts, so that we may speak of fusion action systems or p -local finite group actions without actually referencing the finite group G . For fusion action systems, as is the case with fusion systems, the abstraction process begins with a naive description of a category similar to \mathfrak{X}_G and proceeds by imposing further Saturation Axioms so as to more closely mimic the properties of the ambient case. These Axioms are listed in Definition 2.2.7, then simplified in the results of Subsection 2.5.

With our main object of study defined, we go on to investigate the properties of fusion action systems. These algebraic objects generalize fusion systems, but retain many important aspects of their structure. For instance, it is not difficult to prove a variant of Alperin's Fusion Theorem, so that a saturated fusion action

system is generated by inclusions and the automorphisms of subgroups (Theorem 2.2.12). Topologically, a p -local finite group action arising from an actual finite G -action on X allows us to reconstruct the Borel construction $EG \times_G X$ up to p -completion, just as a p -local finite group arising from G determines the homotopy type of $|BG|_p^\wedge$ (Theorem 3.1.12). Even without reference to an ambient group, there is an obstruction theory to the existence and uniqueness of a p -local finite group action associated to a fusion action system, which mimics that of [BLO2] (Theorem 3.3.5).

The introduction of fusion action systems also leads us to a fusion-theoretic expression of some basic ideas from group theory. If the underlying action of S on X is not faithful, we construct the *core fusion subsystem* that mimics the fusion of the kernel of a nonfaithful ambient group action. We further show that the core subsystem is saturated if the fusion action system is, and that the fusion action system can be viewed as a sort of extension of the core by a finite group (Theorems 2.4.5 and 2.4.13). Moreover, if the fusion action system is transitive, there is a notion of a stabilizer fusion action subsystem of a (well chosen) point, which plays the role of the stabilizer subgroup of that point (Theorem 2.6.4). If we further give ourselves the data of a p -local finite group action associated to our transitive fusion action system, the stabilizer subsystem has an associated p -local finite group whose classifying space is homotopic to the classifying space of the original p -local finite group action (Theorem 3.5.9); this is the fusion-theoretic analogue of the basic fact that if X is a transitive G -set with $H \leq G$ the stabilizer of a point, $EG \times_G X \simeq BH$.

More generally, the purpose of this paper is to develop a framework in which to discuss permutations of finite sets in a fusion-theoretic context. Many of the results are geared toward showing that the structure and technical results of fusion systems or p -local finite groups can be translated to and make sense in our new universe of discourse. Consequently, many of the proofs and much of the development closely mirror the existing literature. We view the possibility of making such a translation as evidence that the definitions of fusion action system and p -local finite group action are the “correct” ones, at least in a moral sense.

0.2 Outline of the paper

Each of this paper’s three sections follow the same basic arc: We are interested certain p -local data that arise from a finite group without reference to that group, so we begin by describing these data when we are actually given a finite group. We then attempt to define what the abstract version of these data should be, only to find that (in the cases of fusion systems and fusion action systems) the naive definition is far too broad to be of much use to us. The solution is to take basic technical properties from the ambient case and impose these as “Saturation Axioms” to define the abstract object of interest. The sections then conclude with further investigations of the properties of these abstract p -local versions of groups and group actions, largely with the goal of connecting our definitions with the existing literature. More specifically:

Section 1 is a review of the notions fusion systems and p -local finite groups, told using the Martino-Priddy conjecture as motivation for the introduction of the more complicated structure of a centric linking system. Although this material is already in the literature (cf. [Lin] and [BLO2] for a more detailed discussion) and seems to be gaining wider recognition, we include our discussion as a prelude to the story we tell about fusion action systems and p -local finite group actions, which we structure along along the same narrative lines.

Section 2 introduces the titular subject of this paper. Subsection 2.1 describes the p -local data we wish to emulate of a finite group G with Sylow subgroup S acting on the finite group X . Subsection 2.2 first gives the abstract definition of a fusion action system, then narrows the terms of discussion by turning the basic results of Subsection 2.1 into the Saturation Axioms for fusion action systems. As a demonstration of the validity of these Axioms, Subsection 2.2 closes with a proof of a simple version of the Alperin Fusion Theorem in the context of fusion action systems.

Fusion action systems arise as an abstraction of the p -local structure of a finite group acting on a finite set, but they turn out to have other interesting interpretations as well. First, a fusion action system \mathfrak{X} has an underlying fusion system $\mathcal{F}^\mathfrak{X}$; this observation allows us to see fusion action systems as generalizations of fusion systems, in the sense that if the set being acted on is trivial, $\mathfrak{X} = \mathcal{F}^\mathfrak{X}$, and moreover we can recover the Saturation Axioms for fusion systems from those of fusion action systems (Subsection 2.2). Second, as the underlying S -action becomes less trivial, a fusion action system looks less like a fusion system and more like the transporter system of a finite group, to the point where if X is a faithful S -set there is a

canonical finite group G that realizes the underlying fusion system as $\mathcal{F}_S(G)$ and the fusion action system \mathfrak{X} as the transporter system of S in G (Subsection 2.3). Finally, every fusion action system has a *core fusion subsystem*, and the fusion action system can be interpreted as an extension of this fusion system by a finite group (Subsection 2.4).

In Subsection 2.5 we address the fact that the Saturation Axioms for fusion action systems are highly redundant by proving that two seemingly weaker sets of Axioms are in fact equivalent. These simplified Saturation Axioms allow us to check saturation with considerably less annoyance, which convenience we use in 2.6. In this final subsection of Section 2, we reconnect fusion action systems with ideas from Puig’s original definition of Frobenius categories in [Pui1]. As motivation, we are able to prove that the stabilizer fusion action subsystem of a fully stabilized point in a transitive saturated fusion action system is itself saturated. In fact, we prove much more, showing in effect that Puig’s theorem on K -normalizers holds in the context of fusion action systems, which will hopefully allow for the discovery of new saturated fusion action systems in the future.

Section 3 turns the focus from pure algebra to algebraic topology by introducing the notion of a classifying space of a fusion action system. The first three subsections mirror the development of p -local finite group theory in [BLO2]: Subsection 3.1 introduces the notion of the X -centric linking action system induced by an ambient group G and shows that this augmentation of the fusion action system contains enough data to reconstruct the homotopy type of the Borel construction $(EG \times_G X)_p^\wedge$. Subsection 3.2 abstractifies this structure so that we are able to talk about X -centric linking action systems associated to the fusion action system \mathfrak{X} without any ambient group. Since it is not at first obvious what the classifying space of a fusion action system should be, we reformulate the notion of linking action system in terms of data contained entirely within the fusion action system itself, namely the associated orbit category $\mathcal{O}^\mathfrak{X}$. The linking action system is shown to provide a solution to a certain homotopy lifting problem indexed on $\mathcal{O}^\mathfrak{X}$. Finally, Subsection 3.3 sets up the obstruction theory to the existence and uniqueness of an X -centric linking action system associated to a given fusion action system, completing the generalization of p -local finite group theory.

In Subsection 3.4 we draw further connections between linking action systems and the existing literature, specifically the abstract transporter systems of [OV]. We show that the former is actually an example of the latter, equipped with the additional data of a homomorphism from the fundamental group of the classifying space to a symmetric group. We further show that this process can be reversed: Given an abstract transporter system \mathcal{T} associated to the fusion system \mathcal{F} and a homomorphism $\theta : \pi_1(|\mathcal{T}|) \rightarrow \Sigma_X$, we can construct a fusion action system \mathfrak{X}^θ with underlying fusion system \mathcal{F} and show that \mathcal{T} contains a linking action system associated to \mathfrak{X}^θ , at least if we assume that \mathcal{T} is defined on enough subgroups. Moreover, we show that \mathfrak{X}^θ is saturated on the subgroups witnessed by \mathcal{T} , suggesting the need for a theorem in the style of [BCG⁺] describing when such a situation would imply that \mathfrak{X}^θ is saturated on all subgroups.

The final Subsection, 3.5, we return to the question of stabilizers of points in transitive fusion action systems introduced in Subsection 2.6, though now we approach the question from the topological perspective of trying to describe the classifying space of the stabilizer. It turns out that, using the results of [OV], the added structure of a linking action systems makes it much easier to prove that the stabilizer fusion subsystem is saturated. Moreover, we get an complete p -local finite group for the stabilizer basically for free, and show that the homotopy type the nerve of the stabilizer centric linking system is the same as the original classifying space for the fusion action system. We should expect this to be true from basic group theory, so perhaps this final result is best interpreted as confirmation that the definitions made in this paper function the way we would hope they do.

1 Background on fusion systems

In this chapter we review the basics of the theory of fusion systems and related concepts. We introduce the classical notion of the fusion system of a finite group and Puig’s abstraction of this idea, the transporter systems of Oliver-Ventura, and the centric linking systems of Broto-Levi-Oliver. As the final portion of this paper will deal with related concepts, we shall also give a brief discussion of the Martino-Priddy Conjecture and its proof by Oliver.

Throughout this paper, let G be a finite group and S a finite p -group, thought of as a Sylow p -subgroup of G when appropriate.

1.1 Transporter and fusion systems of finite groups

We are interested in the p -local structure of finite groups, and furthermore would like to encode these data in a reasonably combinatorial or even category manner. Beyond the one-point category \mathcal{BG} , which by definition has $\mathcal{BG}(*, *) = G$, we may choose to focus on a single Sylow subgroup as an important part of the data.

Definition 1.1.1. Let G be a finite group and $S \in \text{Syl}_p(G)$. The *transporter system on S relative to G* is the category $\mathcal{T}_G = \mathcal{T}_S(G)$ whose objects are all subgroups $P \leq S$ and whose morphisms are given by

$$\mathcal{T}_G(P, Q) = N_G(P, Q) := \{g \in G \mid {}^gP \leq Q\}$$

$N_G(P, Q)$ is the *transporter of P to Q in G* .

This definition singles out a given Sylow subgroup and plays an important role throughout this document. However, it contains too much information, especially p' -data. In fact, \mathcal{T}_G can be easily seen to contain exactly the information of \mathcal{BG} , together with a choice of Sylow subgroup, by noting that $\mathcal{T}_G(1) \cong G$.¹ Indeed, the natural functor $\mathcal{BG} \rightarrow \mathcal{T}_G$ sending $*$ to 1 induces a homotopy inverse to the natural functor $\mathcal{T}_G \rightarrow \mathcal{BG}$; in the world of topology, $|\mathcal{BG}| \simeq |\mathcal{T}_G|$.

One way of understanding the sense in which \mathcal{T}_G has too much information is to note that there may be distinct elements $g, g' \in G$ that conjugate P to Q but are indistinguishable from the point of view of the conjugation action on P . In other words, $c_g|_P = c_{g'}|_P$ or $g^{-1}g' \in Z_G(P)$. Let us suppose that the conjugation action is the truly important p -local data.

Definition 1.1.2. For G a finite group and $S \in \text{Syl}_p(G)$, the *fusion system on S relative to G* is the category $\mathcal{F}_G := \mathcal{F}_S(G)$ whose objects are all subgroups $P \leq S$ and whose morphisms are given by

$$\mathcal{F}_G(P, Q) = \text{Hom}_G(P, Q) := \{\varphi \in \text{Inj}(P, Q) \mid \exists g \in G \text{ s.t. } \varphi = c_g|_P\}$$

Note that we can also write $\mathcal{F}_G(P, Q) = \mathcal{T}_G(P, Q)/Z_G(P)$, so \mathcal{F}_G can be thought of as a quotient of \mathcal{T}_G and we have a natural projection functor $\mathcal{T}_G \rightarrow \mathcal{F}_G$.

Example 1.1.3. The most basic example of a Sylow inclusion $S \leq G$ is the case that the supergroup G is equal to S itself. We denote the resulting fusion system by \mathcal{F}_S , the *minimal fusion system on S* . Minimality in this case means that if H is any finite group with $S \in \text{Syl}_p(H)$ then $\mathcal{F}_S \subseteq \mathcal{F}_S(H) = \mathcal{F}_H$. The importance of this minimal example will become clear with the introduction of abstract fusion systems.

One of the more important properties of fusion systems is that they exhibit “local to global” phenomena. For example, the global condition that $s, s' \in S$ are conjugate in G can be realized by simply looking at automorphism groups in the fusion system itself.

Theorem 1.1.4 (Alperin’s Fusion Theorem). *For G a finite group and $S \in \text{Syl}_p(G)$, the fusion system \mathcal{F}_G has the property that any morphism can be written as a composite of automorphisms in \mathcal{F}_G of subgroups together with subgroup inclusions.*

Proof. See [Alp] for a stronger version of this result. This result was strengthened further in [Gol] to the Alperin-Goldschmidt fusion theorem, which described a particular class of subgroups whose automorphisms determine the fusion system. Finally, Puig showed in [Pui2] that the class of groups identified by Goldschmidt is truly essential in order to generate the fusion system and proved in [Pui1] an abstract analogue of the Fusion Theorem that makes no reference to the finite group G . \square

¹We here make note of our notational convention: Just as for a category \mathcal{C} we denote by $\mathcal{C}(a, b)$ the set of morphisms $\text{Hom}_{\mathcal{C}}(a, b)$, we shall write $\mathcal{C}(a)$ for the automorphism group of the object a .

1.2 Centric linking systems of finite groups

Definition 1.1.2 shows that the fusion system \mathcal{F}_G can be thought of as the quotient of the transporter system \mathcal{T}_G obtained by killing the action of the centralizer of the source. This quotienting process kills both p - and p' -information; if we wish to study the p -local structure of G , perhaps we should seek a less brutal quotient as an intermediary between \mathcal{T}_G and \mathcal{F}_G .

To find this intermediary category, technical considerations suggest that we restrict attention to a particular collection of subgroups of S . The reasons will become clear in short order. Let us therefore introduce a seemingly ad hoc definition of the class of subgroups that will be central in the following discussion:

Definition 1.2.1. A p -subgroup P of G is *p -centric* if $Z(P) \in \text{Syl}_p(Z_G(P))$. Equivalently, P is p -centric if there exists a (necessarily unique) p' -subgroup $Z'_G(P) \leq Z_G(P)$ such that $Z_G(P) = Z(P) \times Z'_G(P)$.

Notation 1.2.2. We shall reserve the notation $Z'_G(P)$ for $O^p(Z_G(P))$ in the case that P is p -centric in G , in which case we also have $Z'_G(P) = O_{p'}(Z_G(P))$.

Notation 1.2.3. By \mathcal{T}_G^c we mean the full subcategory of the transporter system \mathcal{T}_G whose objects are the p -centric subgroups of S . We use similar notation to denote full centric subcategories of fusion systems and other related categorical versions of groups we'll encounter.

We are now in the position to introduce our intermediary between \mathcal{T}_G and \mathcal{F}_G :

Definition 1.2.4. The *centric linking system* of a finite group G with Sylow S is the category \mathcal{L}_G^c whose objects are the p -centric subgroups of S and whose morphisms are the classes

$$\mathcal{L}_G^c(P, Q) = N_G(P, Q)/Z'_G(P)$$

The quotient functors $\mathcal{T}_G^c \longrightarrow \mathcal{L}_G^c \longrightarrow \mathcal{F}_G^c$ relate our three nontrivial notions of G as a category and emphasize how some information is lost at each transition. We shall make use of the relationship between these three players in the sequel.

1.3 The Martino-Priddy Conjecture

The classifying space functor $B : \mathcal{GRP} \rightarrow \mathcal{TOP}$ is the primary tool we use in this document for studying groups in the context of algebraic topology. We can see B as the composition of $\mathcal{B} : \mathcal{GRP} \rightarrow \mathcal{CAT}$ with the geometric realization functor, which suggests that perhaps our alternate categorical versions of finite groups should be viewed as topological spaces via geometric realization.

For the transporter system \mathcal{T}_G , this works: As $|\mathcal{T}_G| \simeq BG$, essentially all the algebraic information of \mathcal{T}_G is realized topologically in this manner. Roughly speaking, the transporter system has a minimal object that includes in every other one. One can see that in such a situation, all the topological data we could hope to extract from the transporter system is in some sense concentrated in the automorphisms of this minimal object; this is a situation that is explored in greater depth in [OV].

However, simply taking the nerve of the fusion system \mathcal{F}_G will not yield an interesting space: The object 1 is initial in \mathcal{F}_G , so $|\mathcal{F}_G|$ is contractible. Indeed, as fusion systems are naturally equipped with a collection of “inclusion morphisms”—honest inclusions of subgroups in this case—this fact could be seen as a special case of the general reason why $|\mathcal{T}_G| \simeq BG$. We will have to be more clever about how we construct a topological space from a fusion system if we are to arrive at anything interesting.

We therefore turn to the centric linking system \mathcal{L}_G^c as an intermediary between transporter systems and fusion systems. The space $|\mathcal{L}_G^c|$ should be related to BG in some way, but as information is lost from the transition from transporter system to linking system it is unreasonable to expect that $|\mathcal{L}_G^c| \simeq BG$.

Example 1.3.1. Let G be your favorite finite group, $S \in \text{Syl}_p(G)$, and H your favorite finite p' -group. Then H is a p' -subgroup of $Z_{G \times H}(S)$ so for any $P \leq S$ we have $H \leq Z'_{G \times H}(P)$. We conclude $\mathcal{L}_G^c \cong \mathcal{L}_{G \times H}^c$ (actual isomorphism of categories). It easily follows that the same result applies to fusion systems: $\mathcal{F}_G = \mathcal{F}_{G \times H}$ (equality of categories).

Since we construct the linking system by killing certain \mathfrak{p}' -primary data of the transporter system, we should look for an operation on topological spaces that “isolates p -information” in some appropriate sense.

Notation 1.3.2. Let $(-)_p^\wedge : \mathcal{TOP} \rightarrow \mathcal{TOP}$ denote the Bousfield-Kan p -completion functor of [BK]. There is a natural transformation $\eta : \text{id}_{\mathcal{TOP}} \Rightarrow (-)_p^\wedge$; for a space \mathcal{X} , let $\eta_{\mathcal{X}} : \mathcal{X} \rightarrow \mathcal{X}_p^\wedge$ denote the resulting p -completion map.

Definition 1.3.3. A space \mathcal{X} is p -complete if the p -completion map $\eta_{\mathcal{X}} : \mathcal{X} \rightarrow \mathcal{X}_p^\wedge$ is a homotopy equivalence. \mathcal{X} is p -good if \mathcal{X}_p^\wedge is p -complete.

The interested reader should refer to [BK] for further properties of the p -completion functor.

So, even though it is unreasonable to expect that $|\mathcal{L}_G^c|$ is homotopy equivalent to BG , it is conceivable that these spaces should be equivalent up to p -completion. The following two theorems of [BLO1] describe the relationship of the centric linking system of G to the space BG_p^\wedge :

Theorem 1.3.4. For G a finite group, the natural functors $BG \longleftarrow \mathcal{T}_G^c \longrightarrow \mathcal{L}_G^c$ induce mod- p cohomology isomorphisms on realization. In particular, $|\mathcal{L}_G^c|_p^\wedge \simeq BG_p^\wedge$.

Theorem 1.3.5 (Weak Martino-Priddy Conjecture). For finite groups G and H , $BG_p^\wedge \simeq BH_p^\wedge$ if and only if the categories \mathcal{L}_G^c and \mathcal{L}_H^c are equivalent.

This still leaves the question of what space we should associate to the fusion system itself. In fact, the surprising result is that the passage from transporter system to fusion system loses no topological information, and so in fact the space associated to \mathcal{F}_G should again be BG_p^\wedge .

Theorem 1.3.6. [Martino-Priddy Conjecture] The finite groups G and H have homotopic p -completed classifying spaces if and only if the p -fusion data of G and H are the same.

Proof. The “topology implies algebra” direction is given in [MP]. The “algebra implies topology” direction was proved by Oliver in [Oli1, Oli2], using the machinery of [BLO2] and the Classification Theorem of Finite Simple Groups. \square

That “the p -fusion data of G and H are the same” means that there is an isomorphism of fusion systems $\mathcal{F}_G \cong \mathcal{F}_H$. The notion of isomorphism of fusion system is much stronger than saying that these categories are equivalent, or even isomorphic as categories: Such a notion would record only the shape of the fusion system as a diagram without giving due deference to the structure of the objects of the fusion system.

Definition 1.3.7. Let G and H be finite groups with respective Sylows S and T . An isomorphism $\alpha : S \rightarrow T$ is a *fusion preserving isomorphism* if for every $P, Q \leq S$ and $\beta \in \text{Hom}(P, Q)$, $\beta \in \text{Hom}_G(P, Q)$ if and only if $\alpha\beta\alpha^{-1} \in \text{Hom}_H(\alpha P, \alpha Q)$. In this case the fusion systems \mathcal{F}_G and \mathcal{F}_H are *isomorphic as fusion systems*.

So we claim that the data of the fusion system \mathcal{F}_G determine the p -completed homotopy type of BG , which returns us to the question of exactly how to associate a topological space to a fusion system. We have already seen that simply taking the nerve of \mathcal{F}_G yields nothing interesting. This should not be surprising, as taking the geometric realization of the fusion system only records the shape of the category as a diagram without taking into account the fact that it is a diagram *in p -groups*. This is not a problem for either the transporter or linking systems of G , as for any $P \leq S$ we have $P \leq N_G(P)$ and therefore there is a natural way to identify P with a subgroup of its automorphism group. Such is not the case for fusion systems, so we must try a little harder to recover this information. In so doing, we introduce yet another categorical version of the finite group G .

Definition 1.3.8. The *orbit category*² of \mathcal{F}_G is the category $\mathcal{O}_G := \mathcal{O}(\mathcal{F}_G)$ whose objects are the subgroups of S and whose morphisms are given by

$$\mathcal{O}_G(P, Q) = Q \backslash \mathcal{F}_G(P, Q)$$

²This notion is not to be confused with the category of orbits of G , whose objects are the transitive G -sets and where morphisms are maps of G -sets. Although there is a relationship between these two notions, we will not make use of the category of G -orbits in this document.

In other words, the hom-set from P to Q is the set orbits of the Q -action of $\mathcal{F}_G(P, Q)$ given by postcomposition by c_q .

\mathcal{O}_G^c will denote the full subcategory of \mathcal{O}_G whose objects are the \mathcal{F} -centric subgroups of S .

The functor $B- : \mathcal{F}_G \rightarrow \mathcal{TOP}$ does not descend to a functor $\mathcal{O}_G \rightarrow \mathcal{TOP}$, but because \mathcal{O}_G is defined by quotienting out inner automorphisms, it is easy to see that there is a homotopy functor $\overline{B}- : \mathcal{O}_G \rightarrow ho\mathcal{TOP}$. If we could find a homotopy lifting

$$\begin{array}{ccc} & \mathcal{TOP} & \\ \tilde{B}- \nearrow & \downarrow & \\ \mathcal{O}_G & \xrightarrow{\overline{B}-} & ho\mathcal{TOP} \end{array}$$

we could consider $\text{hocolim}_{\mathcal{O}_G} \tilde{B}-$, and relate this space to BG_p^\wedge . The following Proposition explains this relationship.

Proposition 1.3.9. *For G a finite group, consider the diagram*

$$\begin{array}{ccc} \mathcal{L}_G^c & \xrightarrow{*} & \mathcal{TOP} \\ \pi \downarrow & \nearrow L & \\ \mathcal{O}_G^c & & \end{array}$$

where π is the composite of the natural quotients $\mathcal{L}_G^c \rightarrow \mathcal{F}_G^c \rightarrow \mathcal{O}_G^c$ and L is the left homotopy Kan extension of the trivial functor $*$ over π . Then L is a homotopy lifting of $\overline{B}- : \mathcal{O}_G^c \rightarrow ho\mathcal{TOP}$, and in particular we have

$$\text{hocolim}_{\mathcal{O}_G^c} L \simeq \text{hocolim}_{\mathcal{L}_G^c} * = |\mathcal{L}_G^c| \simeq_p BG$$

Proof. [BLO2]. □

So far we have simply restated the original question of whether topological information is lost on the transition from linking system to fusion system: If we have a linking system in mind for \mathcal{F}_G , there is a homotopy lifting of $\overline{B}-$, which allows us to construct our desired space from the fusion system. But what if there is another finite group H such that $\mathcal{F}_G = \mathcal{F}_H$ and yet $\mathcal{L}_G^c \neq \mathcal{L}_H^c$: Is it possible there are two distinct homotopy liftings and thus two different spaces associated to \mathcal{F}_G ? Or can this never happen? How do we approach this problem?

1.4 Group theory without groups

The basic problem introduced at the end of Subsection 1.3 is the need to think of fusion and linking systems as algebraic objects distinct from the finite groups from which they came.

Puig provided the necessary insight and abstraction to codify this generalization. Here we introduce his idea of abstract fusion systems (or “Frobenius categories” in the terminology of [Pui1]), though we shall use the language of Broto-Levi-Oliver. We also review the abstraction of the notion of centric linking system, due to [BLO2].

Definition 1.4.1. Let S be a p -group. An *abstract fusion system on S* is a category \mathcal{F} whose objects are all subgroups $P \leq S$ and whose morphisms are some collection of injective group maps: $\mathcal{F}(P, Q) \subseteq \text{Inj}(P, Q)$. We require that the following conditions be satisfied:

- (**S -conjugacy**) The minimal fusion system \mathcal{F}_S is a subcategory of \mathcal{F} .
- (**Divisibility**) Every morphism of \mathcal{F} factors as an isomorphism of groups followed by an inclusion.

Composition of morphisms is composition of group maps.

This is a very simple definition. In fact, it is perhaps too simple to be useful: This mimics the situation where S is a p -subgroup of some unnamed ambient group, but not where S is a *Sylow* p -subgroup. There is a great deal of additional structure that comes from such a Sylow inclusion; the question is how to codify these interesting data without reference to an ambient group. This will lead to the addition of *Saturation Axioms* that must be imposed on a fusion system.

Definition 1.4.2. We will need the following terms to state the Saturation Axioms:

- $P \leq S$ is *fully normalized* in \mathcal{F} if $|N_S(P)| \geq |N_S(Q)|$ for all $Q \cong_{\mathcal{F}} P$.
- $P \leq S$ is *fully centralized* in \mathcal{F} if $|Z_S(P)| \geq |Z_S(Q)|$ for all $Q \cong_{\mathcal{F}} P$.
- For any $\varphi \in \mathcal{F}(P, Q)_{\text{iso}}$, let $N_{\varphi} \leq N_S(P)$ denote the group

$$\begin{aligned} N_{\varphi} &= \{n \in N_S(P) \mid \varphi \circ c_n \circ \varphi^{-1} \in \text{Aut}_S(Q)\} \\ &= \{n \in N_S(P) \mid \exists s \in S \text{ s.t. } \forall p \in P, \varphi({}^n p) = {}^s \varphi(p)\} \end{aligned}$$

N_{φ} will be called the *extender* of φ .

Remark 1.4.3. Perhaps some motivation for these concepts is in order. Each of these definitions comes from the idea that there are certain “global” phenomena that can be captured purely through local, fusion-theoretic data of a group. For instance, if there is an ambient Sylow G giving rise to the fusion system, then $P \leq S$ is fully normalized in \mathcal{F} if and only if $N_S(P) \in \text{Syl}_p(N_G(P))$, and similarly for the concept of full centralization.

The motivation for the extender N_{φ} comes from Alperin’s Fusion Theorem 1.1.4. If we wish that the morphisms of a fusion system be generated by inclusions and automorphisms of subgroups, there must be some way of extending certain morphisms between different subgroups within the fusion system. The extender is the maximal subgroup of $N_S(P)$ to which we could hope to extend $\varphi \in \mathcal{F}(P, Q)$, so the question becomes when we can achieve this maximal extension.

Definition 1.4.4 (Saturation Axioms for Abstract Fusion Systems). The fusion system \mathcal{F} is *saturated* if

- Whenever P is fully \mathcal{F} -normalized, P is fully \mathcal{F} -centralized.
- Whenever P is fully \mathcal{F} -normalized, $\text{Aut}_S(P) \in \text{Syl}_p(\mathcal{F}(P))$.
- If Q is fully \mathcal{F} -centralized and $\varphi \in \text{Iso}_{\mathcal{F}}(P, Q)$, then there is some morphism $\tilde{\varphi} \in \mathcal{F}(N_{\varphi}, S)$ that extends φ : $\tilde{\varphi}|_P = \varphi$.

The first two conditions will be referred to as the *Sylow Axioms* and the third the *Extension Axiom*.

Note that the Saturation Axioms are somewhat redundant, in that seemingly less restrictive Axioms turn out to be equivalent. Examples of such simplified axiom sets can be found in [OS] and [RS]. See also [Pui1] for a different take on the nature of saturation.

Saturated fusion systems form a category, though we shall not make great use of the notion of morphism of fusion system in this paper.

Not only can we discuss fusion systems without reference to an ambient group, but there also exists a notion of abstract transporter systems that generalizes both the constructions \mathcal{T}_G and \mathcal{L}_G^c .

Definition 1.4.5 ([OV]). Let \mathcal{F} be an abstract fusion system on S . An *abstract transporter system associated to \mathcal{F}* is a category \mathcal{T} whose objects are some set of subgroups $P \leq S$ that is closed under \mathcal{F} -conjugacy and overgroups, together with functors

$$\mathcal{T}_S^{\text{Ob}(\mathcal{T})}(S) \xrightarrow{\delta} \mathcal{T} \xrightarrow{\pi} \mathcal{F}$$

For any $s \in N_S(P, Q)$, set $\hat{s} = \delta_{P,Q}(s) \in \mathcal{T}(P, Q)$. Similarly, for any $\mathfrak{g} \in \mathcal{T}(P, Q)$, set $c_{\mathfrak{g}} = \pi_{P,Q}(\mathfrak{g}) \in \mathcal{F}(P, Q)$.

The following Axioms apply:

(A1) On objects, δ is the identity and π is the inclusion.

(A2) For any $P \in \text{Ob}(\mathcal{T})$, define

$$E(P) = \ker [\pi_{P,P} : \mathcal{T}(P) \rightarrow \mathcal{F}(P)]$$

Then for any $P, Q \in \text{Ob}(\mathcal{T})$, the group $E(P)$ acts right-freely and $E(Q)$ acts left-freely on $\mathcal{T}(P, Q)$. Moreover, the map $\pi_{P,Q} : \mathcal{T}(P, Q) \rightarrow \mathcal{F}(P, Q)$ is the orbit map of the $E(P)$ -action.

(B) The functor δ is injective on morphisms, and for all $s \in N_S(P, Q)$ we have $c_s = c_s \in \mathcal{F}(P, Q)$.

(C) For all $\mathfrak{g} \in \mathcal{T}(P, Q)$ and all $p \in P$, the following diagram commutes in \mathcal{T} :

$$\begin{array}{ccc} P & \xrightarrow{\mathfrak{g}} & Q \\ \widehat{p} \downarrow & & \downarrow \widehat{c_{\mathfrak{g}}(p)} \\ Q & \xrightarrow{\mathfrak{g}} & Q \end{array}$$

Saturation axioms:

(I) $\delta_{S,S}(S) \in \text{Syl}_p(\mathcal{T}(S))$.

(II) For all $\mathfrak{g} \in \text{Iso}_{\mathcal{T}}(P, Q)$ and supergroups $\tilde{P} \supseteq P$ and $\tilde{Q} \supseteq Q$ such that $\mathfrak{g} \circ \delta_{P,P}(\tilde{P}) \circ \mathfrak{g}^{-1} \leq \delta_{Q,Q}(\tilde{Q})$, there is a morphism $\tilde{\mathfrak{g}} \in \mathcal{T}(\tilde{P}, \tilde{Q})$ that satisfies $\tilde{\mathfrak{g}} \circ \widehat{1}_{\tilde{P}} = \widehat{1}_{\tilde{Q}} \circ \mathfrak{g} \in \mathcal{T}(P, Q)$.

Notation 1.4.6. As was noted in the definition of transporter systems, for any $s \in N_S(P, Q)$ we set

$$\widehat{s}|_P^Q = \delta_{P,Q}(s) \in \mathcal{T}(P, Q)$$

and if the source and target are obvious from the context we shall simply write \widehat{s} . Similarly, for any $R \leq N_S(P)$ we denote by $\widehat{R}|_P^P \leq \mathcal{T}(P)$ the group $\delta_{P,P}(R)$, again writing simply \widehat{R} if there is no chance of confusion.

Remark 1.4.7. $E(P)$ always contains $\delta_{P,P}(Z(P))$ by Axiom (B). As δ is injective, we identify $Z(P)$ with its image in $\mathcal{T}(P)$.

We say that $P \leq S$ is \mathcal{F} -centric if $Z_S(Q) = Z(Q)$ for all Q that are \mathcal{F} -conjugate to P . The collection of \mathcal{F} -centric subgroups is of central importance to the study of the homotopy theory of fusion systems.

Definition 1.4.8. Let \mathcal{F} be a saturated fusion system on S . A transporter system \mathcal{L} associated to \mathcal{F} is an *abstract centric linking system* if:

- $\text{Ob}(\mathcal{L})$ is the collection of \mathcal{F} -centric subgroups of S .
- For every $P \in \text{Ob}(\mathcal{L})$, $E(P) = Z(P)$.

Thus an abstract linking system can be thought of as a minimal transporter system on the \mathcal{F} -centric subgroups of S .

The rationale for augmenting a fusion system with the more complicated structure of a centric linking system is that they lead to a proof of the Martino-Priddy Conjecture: We already know that the centric linking system \mathcal{L}_G^c arising from a finite group G determines the homotopy type of BG_p^\wedge , and using the obstruction theory of [BLO2] and the Classification of Finite Simple Groups, Oliver was able to show in [Oli1, Oli2] that there is a unique centric linking system associated to any fusion system of a finite group.

2 Fusion action systems

2.1 The ambient case

To motivate the following discussion, let us identify p -local data of a finite group acting on a finite set that we are trying to abstract. Let G be a finite group, X a G -set, and $S \in \text{Syl}_p(G)$. We use the shorthand $\mathcal{F}_G = \mathcal{F}_S(G)$ and $\mathcal{F}_S = \mathcal{F}_S(S)$. We additionally denote the action map defining the G -set structure on X by ℓ , and will write ℓ_g or ℓ_H for the permutation of X represented by g or the subgroup of Σ_X that is the image of $H \leq G$, respectively.

Definition 2.1.1. The *fusion action system on S relative to the G -action on X* is the category $\mathfrak{X}_G := \mathfrak{X}_S(G)$ whose objects are the subgroups $P \leq S$ and whose morphisms are given by

$$\mathfrak{X}_G(P, Q) = \{(\varphi, \sigma) \in \text{Hom}(P, Q) \times \Sigma_X \mid \exists g \in G \text{ s.t. } \varphi = c_g|_P \in \text{Inj}(P, Q) \text{ and } \sigma = \ell_g : {}_P X \cong {}_P X\}.$$

Any pair (φ, σ) that appears as a morphism in $\mathfrak{X}_G(P, Q)$ has the property that for all $p \in P$ and $x \in X$, $\sigma(p \cdot x) = \varphi(p) \cdot \sigma(x)$, or equivalently that $\sigma \ell_p \sigma^{-1} = \ell_{\varphi(p)}$; such a pair of an injection of subgroups and permutation of X is *intertwined*.

Example 2.1.2. The restriction of the G -action to S defines a fusion action system \mathfrak{X}_S , the *minimal fusion action system* of a given S -action on X .

The fusion action system \mathfrak{X}_G has an *underlying fusion system* on S , obtained by simply ignoring the second coordinates of the morphisms of \mathfrak{X}_G . This is of course \mathcal{F}_G .

If we ignore the first coordinates of \mathfrak{X}_G , we obtain another interesting algebraic structure, which turns out to be, effectively, a transporter system on a quotient of S . Then $\ell_S \leq \text{Syl}_p(\ell_G)$, and we can talk about the transporter system on ℓ_S relative to this inclusion, $\mathcal{T}_{\ell_G} := \mathcal{T}_{\ell_S}(\ell_G)$. Then the functor $\mathfrak{X}_G \rightarrow \mathcal{T}_{\ell_G}$ given on objects by $P \mapsto \ell_P$ and on morphisms by projection onto the second factor is surjective, in the sense that any morphism of \mathcal{T}_{ℓ_G} lies in the image of a morphism of \mathfrak{X}_G .

Definition 2.1.3. For any pair of subgroups $P, Q \leq S$, let $\pi_{\Sigma}^{P, Q} : \mathfrak{X}_G(P, Q) \rightarrow \Sigma_X$ be the set-map projection onto the second coordinate. We shall suppress the reference to the source and target groups and simply write π_{Σ} for this map. When $P = Q$, the set-map π_{Σ} is a homomorphism of groups. In this case let $\Sigma_{\mathfrak{X}}^G(P)$ be the image of $\mathfrak{X}_G(P)$ under π_{Σ} , and similarly set $\Sigma_{\mathfrak{X}}^S(P)$ to be the image of $\mathfrak{X}_S(P)$ under π_{Σ} .

To emphasize the connection with fusion systems, we introduce the following notation:

Notation 2.1.4. Let $\text{Aut}_G(P; X)$ denote the group $\mathfrak{X}_G(P)$. This group is a simultaneous generalization of $\text{Aut}_G(P) = \mathcal{F}_G(P)$ and the group $\text{Aut}_G(X)$ of G -set automorphisms of X . Similarly define $\text{Aut}_S(P; X)$.

There are several groups of automorphisms that arise from inspection of the category \mathfrak{X}_G , two of which are $\text{Aut}_G(P; X)$ and $\text{Aut}_S(P; X)$. We need some notation to introduce the others.

Notation 2.1.5. Let \widehat{C} be the *core*, or kernel, of the G action on X . We define the *X -normalizer* and *X -centralizer in G* of a subgroup $H \leq G$ to be

$$N_G(H; X) := N_G(H) \cap \widehat{C} \quad \text{and} \quad Z_G(H; X) := Z_G(H) \cap \widehat{C}.$$

Similarly, let C be the core of the S -action on ${}_S X$, so $C = \widehat{C} \cap S$. The *X -normalizer* and *X -centralizer in S* of $P \leq S$ are then

$$N_S(P; X) := N_S(P) \cap C \quad \text{and} \quad Z_S(P; X) := Z_S(P) \cap C.$$

Note that $N_G(H; X)$ is just another name for $N_{\widehat{C}}(H)$. We use this notation to emphasize the idea that G is acting simultaneously on its subgroups (by conjugation) and on X (by left multiplication).




Figure 1: Naming the G -automizer groups

Definition 2.1.6. For any $P \leq S$, we have the inclusions of the groups $Z_G(P; X)$, $N_G(P; X)$, $Z_G(P)$, and $N_G(P)$ as depicted in Figure 1. All of these inclusions are normal, so we name to the respective quotients in that Figure as well. The term G -automizers of P will refer to any of these quotient groups.

We can expand on the definitions of these G -automizers as follows:

$$\begin{aligned}
\mathfrak{X}_G(P) &= N_G(P)/Z_G(P; X) &= \text{Aut}_G(P; X) \\
\mathcal{F}_G(P) &= N_G(P)/Z_G(P) &= \{\varphi \in \text{Aut}(P) \mid (\varphi, \sigma) \in \mathfrak{X}_G(P)\} \\
\Sigma_{\mathfrak{X}}^G(P) &= N_G(P)/N_G(P; X) &= \{\sigma \in \Sigma_X \mid (\varphi, \sigma) \in \mathfrak{X}_G(P)\} \\
\mathcal{F}_G(P)_0 &= N_G(P; X)/Z_G(P; X) &= \{\varphi \in \text{Aut}(P) \mid (\varphi, \text{id}_X) \in \mathfrak{X}_G(P)\} \\
\Sigma_{\mathfrak{X}}^G(P)_0 &= Z_G(P)/Z_G(P; X) &= \{\sigma \in \Sigma_X \mid (\text{id}_P, \sigma) \in \mathfrak{X}_G(P)\}
\end{aligned}$$

We also have the short exact sequences

$$\begin{aligned}
1 &\longrightarrow \Sigma_{\mathfrak{X}}^G(P)_0 \longrightarrow \mathfrak{X}_G(P) \longrightarrow \mathcal{F}_G(P) \longrightarrow 1, \\
1 &\longrightarrow \mathcal{F}_G(P)_0 \longrightarrow \mathfrak{X}_G(P) \longrightarrow \Sigma_{\mathfrak{X}}^G(P) \longrightarrow 1.
\end{aligned}$$

Clearly $\Sigma_{\mathfrak{X}}^G(P)_0$ can be identified with a subgroup of $\text{Aut}_P(X)$ and $\mathcal{F}_G(P)_0$ with a subgroup of $\mathcal{F}_G(P)$. If $\varphi \in \mathcal{F}_G(P)_0$, the identity map defines an isomorphism of P -sets $\text{id}_X : {}_P X \cong {}_P X$, or $\ell_p = \ell_{\varphi(p)}$ for all $p \in P$. Thus $\varphi(p) = p \bmod \widehat{C}$, so we have $\mathcal{F}_G(P)_0 \leq \ker(\mathcal{F}_G(P) \rightarrow \mathcal{F}_{\ell_G}(\ell_P))$.

Relative to S we have the same relationships amongst the groups $Z_S(P; X)$, $Z_S(P)$, $N_S(P; X)$, and $N_S(P)$. We also have the inclusions $Z_S(P; X) \leq Z_G(P; X)$, etc. These data give rise to the rather more complicated diagram of Figure 2.

This diagram allows us to relate the “minimal” automizer groups—those that arise from the simultaneous S -action on its subgroups and X —with those that arise from G . For certain subgroups $P \leq S$, there is a particularly nice relationship; to express it, we will need the following terminology:

Definition 2.1.7. Given the fusion action system \mathfrak{X}_G and $P \leq S$, we say that

- P is *fully normalized relative to G* if $N_S(P) \in \text{Syl}_p(N_G(P))$.
- P is *fully centralized relative to G* if $Z_S(P) \in \text{Syl}_p(Z_G(P))$.
- P is *fully X -normalized relative to G* if $N_S(P; X) \in \text{Syl}_p(N_G(P; X))$.
- P is *fully X -centralized relative to G* if $Z_S(P; X) \in \text{Syl}_p(Z_G(P; X))$.




Figure 2: Comparing S -automizers to G -automizers

The reference to G will be omitted if it is clear from the context.

Just as in the case of ordinary fusion systems, we can think of these terms in the following manner: Instead of starting out with a chosen Sylow subgroup $S \leq G$ and a p -subgroup $P \leq S$, we want to investigate the p -subgroup P on its own terms. For instance, to understand the p -part of $N_G(P)$ we must pick a “right” Sylow of G : Such a Sylow must contain not just P , but also a Sylow subgroup of $N_G(P)$. Saying that P is fully normalized means that we have made this choice correctly within the G -conjugacy class of P inside S . Moreover, Sylow’s theorems show that for a fixed Sylow S and $P \leq S$, there is a subgroup $Q \leq S$ such that Q is G -conjugate to P and Q is fully normalized, etc.

We now use this terminology to describe the relationships between the S - and G -automizers of P .

Proposition 2.1.8. *Fix \mathfrak{X}_G and $P \leq S$.*

(1) *If P is fully normalized relative to G , then*

- $\text{Aut}_S(P) \in \text{Syl}_p(\mathcal{F}_G(P))$
- $\text{Aut}_S(P; X) \in \text{Syl}_p(\mathfrak{X}_G(P))$
- $\Sigma_{\mathfrak{X}}^S(P) \in \text{Syl}_p(\Sigma_{\mathfrak{X}}^G(P))$

and furthermore P is fully centralized, X -normalized, and X -centralized relative to G .

(2) *If P is fully X -normalized relative to G , then $\mathcal{F}_S(P)_0 \in \text{Syl}_p(\mathcal{F}_G(P)_0)$ and P is fully X -centralized relative to G .*

(3) If P is fully centralized relative to G , then $\Sigma_{\mathfrak{X}}^S(P)_0 \in \text{Syl}_p(\Sigma_{\mathfrak{X}}^G(P)_0)$ and P is fully X -centralized relative to G .

Proof. If P is fully normalized, consider the diagram of short exact sequences:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & Z_G(P; X) & \longrightarrow & N_G(P; X) & \longrightarrow & \mathfrak{X}_G(P) & \longrightarrow & 1 \\ & & \uparrow & & \uparrow \text{SyLOW} & & \uparrow & & \\ 1 & \longrightarrow & Z_S(P; X) & \longrightarrow & N_S(P; X) & \longrightarrow & \text{Aut}_S(P; X) & \longrightarrow & 1 \end{array}$$

The middle inclusion's being SyLOW implies that the other two are as well, and thus $\text{Aut}_S(P; X) \in \text{Syl}_p(\mathfrak{X}_G(P))$ and P is fully X -centralized. The other claims are proved in a similar manner. \square

Lemma 2.1.8 tells us certain properties of \mathfrak{X}_G that follow from a subgroup's having the property that the normalizer, X -normalizer, or centralizer of that subgroup in S is SyLOW in the respective group of G . It gives no information, however, about the case when P is fully X -centralized relative to G , which turns out to be very important for understanding extensions of morphisms in the fusion action system. Recall that if $g \in G$ is such that ${}^gP \leq S$, then g determines the morphism $(c_g, \ell_g) \in \mathfrak{X}_G(P, {}^gP)$.

Definition 2.1.9. For $P \leq S$ and $g \in G$ such that ${}^gP \leq S$, define

$$N_{(c_g, \ell_g)} = \{n \in N_S(P) \mid \exists s \in N_S({}^gP) \text{ s.t. } (c_{gng^{-1}}, \ell_{gng^{-1}}) = (c_s, \ell_s)\}.$$

The group $N_{(c_g, \ell_g)}$ is the *extender* of (c_g, ℓ_g) in the fusion action system.

Remark 2.1.10. Consider the following easy observations:

- $P \cdot Z_S(P; X) \leq N_{(c_g, \ell_g)} \leq N_S(P)$.
- We define a notion of “translation” along morphisms: Given $g \in G$ and $P \leq S$ such that ${}^gP \leq S$, set

$$t_{(c_g, \ell_g)} : \mathfrak{X}_G(P) \rightarrow \mathfrak{X}_G({}^gP) : (\varphi, \sigma) \mapsto (c_g \varphi c_g^{-1}, \ell_g \sigma \ell_g^{-1}).$$

We could then define $N_{(c_g, \ell_g)}$ to be the subgroup of $n \in N_S(P)$ such that

$$(c_n, \ell_n) \in t_{(c_g, \ell_g)}^{-1}(\text{Aut}_S({}^gP; X))$$

or even more confusingly, the preimage in $N_S(P)$ of the preimage in $\mathfrak{X}_G(P)$ of $\text{Aut}_S({}^gP; X)$.

- The extender $N_{(c_n, \ell_n)}$ is the largest subgroup of $N_S(P)$ to which we could hope to extend (c_g, ℓ_g) . Here, extension of a morphism means that we would find a $g' \in G$ such that ${}^{g'}N_{(c_g, \ell_g)} \leq S$, $\ell_{g'} = \ell_g$, and $c_{g'} = c_g$ on P . The last two conditions are equivalent to requiring that $g'Z_G(P; X) = gZ_G(P; X)$. In this case we say that $(c_{g'}, \ell_{g'}) = (c_{g'}, \ell_g) \in \mathfrak{X}_G(N_{(c_g, \ell_g)}, S)$ is an extension of $(c_g, \ell_g) \in \mathfrak{X}_G(P, S)$.

The reason why $N_{(c_g, \ell_g)}$ is the domain of the largest possible extension is as follows: Pick some $n \in N_S(P)$ and imagine there some $g' \in G$ such that ${}^{g'}\langle P, n \rangle \leq S$ and $(c_g, \ell_g) = (c_{g'}|_P, \ell_{g'})$. Then ${}^{g'}n \in N_S({}^gP)$, so there is some $s \in N_S({}^gP)$ such that $(c_{g'n(g')^{-1}}, \ell_{g'n(g')^{-1}}) = (c_s, \ell_s)$. On the other hand, the assumption that $c_g = c_{g'}$ on P implies that, on restriction, we have $c_{g'n(g')^{-1}} = c_{gng^{-1}}$, and similarly we have $\ell_{g'n(g')^{-1}} = \ell_{gng^{-1}}$, which tells us that $n \in N_{(c_g, \ell_g)}$, as desired.

Lemma 2.1.11. Let $P \leq S$ and $g \in G$ be such that ${}^gP \leq S$ is fully X -centralized in G . Then there is a $g' \in G$ such that ${}^{g'}N_{(c_g, \ell_g)} \leq S$ and $(c_{g'}|_P, \ell_{g'}) = (c_g, \ell_g)$. In other words, g' defines an extension of (c_g, ℓ_g) to a morphism in $\mathfrak{X}_G(N_{(c_g, \ell_g)}, S)$.

Proof. If gP is fully X -centralized, then

$$\begin{aligned} [N_S({}^gP) : N_S({}^gP) \cdot Z_G({}^gP; X)] &= \frac{|N_S({}^gP)| |Z_G({}^gP; X)| / |Z_S({}^gP; X)|}{|N_S({}^gP)|} \\ &= [Z_S({}^gP; X) : Z_G({}^gP; X)] \end{aligned}$$

is prime to p , so $N_S({}^gP) \in \text{Syl}_p(N_S({}^gP) \cdot Z_G({}^gP; X))$. From the definition of the extender we have ${}^gN_{(c_g, \ell_g)}$ is a p -subgroup of $N_S({}^gP) \cdot Z_G({}^gP; X)$, so we can choose some $z \in Z_G({}^gP; X)$ so that ${}^{zg}N_{(c_g, \ell_g)} \leq N_S({}^gP)$. Then setting $g' = zg$ gives the desired extension $(c_{g'}, \ell_{g'})$. \square

2.2 Abstract fusion action systems

In this section we describe an abstraction of the fusion action systems of Section 2.1 that makes no reference to an ambient group. Let us start by naming our universe of discourse:

Definition 2.2.1. Let $\mathfrak{U} := \mathfrak{U}(S; X)$ be the category whose objects are all subgroups $P \leq S$ and whose morphisms are given by

$$\mathfrak{U}(P, Q) = \{(\varphi, \sigma) \mid \varphi \in \text{Inj}(P, Q), \sigma \in \Sigma_X, \text{ and } \varphi \text{ intertwines } \sigma\}.$$

Recall that the condition on the pair (φ, σ) means that for all $p \in P$ we have $\sigma \ell_p \sigma^{-1} = \ell_{\varphi(p)}$.

By identifying \mathfrak{U} as the “universe of discourse,” we mean that this is the category in which all of our abstract fusion action systems will live.

Definition 2.2.2. An *abstract fusion action system* of S acting on X is a category \mathfrak{X} such that $\mathfrak{X}_S \subseteq \mathfrak{X} \subseteq \mathfrak{U}$ and that satisfies the *Divisibility Axiom*: Every morphism of \mathfrak{X} factors as an isomorphism followed by an inclusion.

The *underlying fusion system* of \mathfrak{X} is the fusion system $\mathcal{F}^{\mathfrak{X}}$ on S , where

$$\mathcal{F}^{\mathfrak{X}}(P, Q) = \{\varphi \in \text{Inj}(P, Q) \mid \exists \sigma \in \Sigma_X \text{ s.t. } (\varphi, \sigma) \in \mathfrak{X}(P, Q)\}.$$

Thus $\mathcal{F}^{\mathfrak{X}}$ is the fusion system obtained by projecting the morphisms of \mathfrak{X} onto the first coordinate.

By referring to this underlying fusion system we are able to speak of $P \leq S$ being fully normalized or fully centralized (by definition, in $\mathcal{F}^{\mathfrak{X}}$), as well as fully X -normalized and X -centralized. The latter two are again formed by considering the normalizers or centralizers of subgroups intersected with C , the core of the action of S on X .

Lemma 2.2.3. *The core C is strongly closed in $\mathcal{F}^{\mathfrak{X}}$, i.e., if φ is a morphism of $\mathcal{F}^{\mathfrak{X}}$ and $c \in C$ is in the source of φ , then $\varphi(c) \in C$.*

Proof. If $c \in C$ and $\varphi \in \mathcal{F}^{\mathfrak{X}}(\langle c \rangle, S)$, then there is some $(\varphi, \sigma) \in \mathfrak{X}(\langle c \rangle, S)$. Then φ and σ 's being intertwined implies that $\ell_{\varphi(c)} = \sigma \ell_c \sigma^{-1} = \text{id}_X$ and thus $\varphi(c) \in c$. \square

Of course, we need further axioms to restrict ourselves to the interesting cases. The goal is to turn the observations of Lemmas 2.1.8 and 2.1.11 into the *Saturation Axioms* for abstract fusion action systems.

Note that all the terms of Definition 2.1.6 have obvious analogues in this context, so we may talk about the group the projection maps $\pi_{\mathcal{F}} : \mathfrak{X}(P, Q) \rightarrow \mathcal{F}^{\mathfrak{X}}(P, Q)$ and $\pi_{\Sigma} : \mathfrak{X}(P, Q) \rightarrow \Sigma_X$. In particular, we have $\Sigma_{\mathfrak{X}}(P)$ is the group of permutations of X that appear as a second coordinate of some automorphism of P in \mathfrak{X} , while $\mathcal{F}^{\mathfrak{X}}(P)_0$ and $\Sigma_{\mathfrak{X}}(P)_0$ are the kernels of the natural projections $\mathfrak{X}(P) \rightarrow \Sigma_{\mathfrak{X}}(P)$ and $\mathfrak{X}(P) \rightarrow \mathcal{F}^{\mathfrak{X}}(P)$, respectively. All of these groups will be referred to as the \mathfrak{X} -*automizers* of the subgroup P .

We thus find ourselves in the situation where'd like to compare the various \mathfrak{X} -automizers to their S -automizer subgroups, and require that these inclusions be Sylow in certain circumstances. Unfortunately, without reference to an ambient group G , we do not have the entirety of Figure 2, but only the truncated Figure 3.

What we need is an analogue of P 's being fully normalized, etc., that makes no reference to an ambient group.

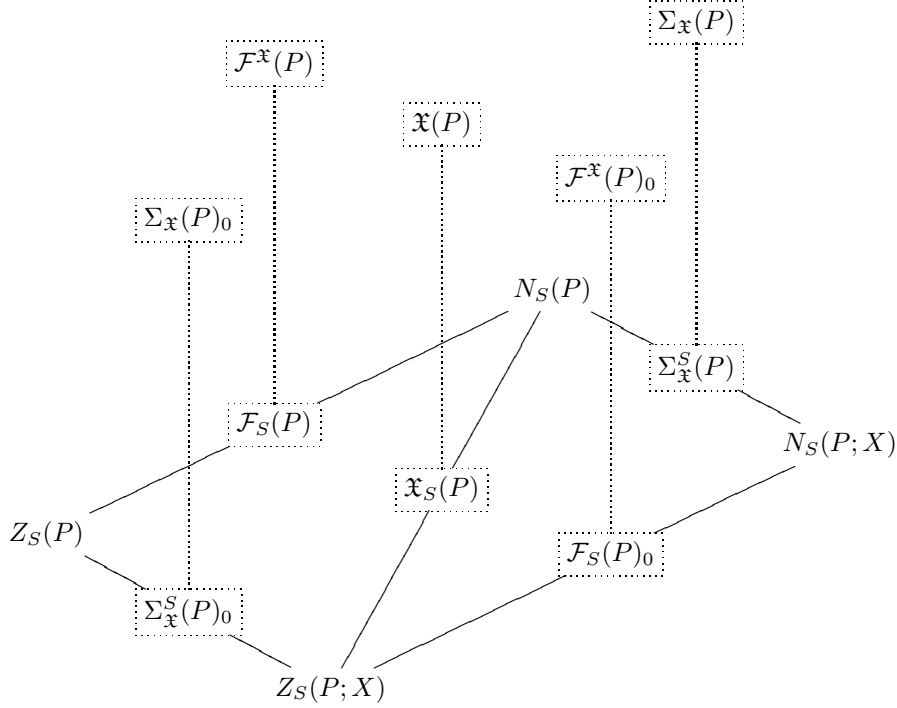


Figure 3: Comparing S -automorphisms to \mathfrak{X} -automorphisms

Definition 2.2.4. Fix a fusion action system \mathfrak{X} and a subgroup $P \leq Q$

- P is *fully normalized in \mathfrak{X}* if $|N_S(P)| \geq |N_S(Q)|$ for all $Q \cong_{\mathcal{F}^{\mathfrak{X}}} P$.
- P is *fully centralized in \mathfrak{X}* if $|Z_S(P)| \geq |Z_S(Q)|$ for all $Q \cong_{\mathcal{F}^{\mathfrak{X}}} P$.
- P is *fully X -normalized in \mathfrak{X}* if $|N_S(P; X)| \geq |N_S(Q : X)|$ for all $Q \cong_{\mathcal{F}^{\mathfrak{X}}} P$.
- P is *fully X -centralized in \mathfrak{X}* if $|Z_S(P; X)| \geq |Z_S(Q : X)|$ for all $Q \cong_{\mathcal{F}^{\mathfrak{X}}} P$.

If $\mathfrak{X} = \mathfrak{X}_G$ for some finite group G , it is well known (cf. [BLO2], Proposition 1.3) that P is fully normalized in $\mathcal{F}^{\mathfrak{X}} = \mathcal{F}_G$ if and only if P is fully normalized with respect to G , and similarly for both definitions of full centralization. The following Lemma shows that the same is true for the two new terms we have introduced

Lemma 2.2.5. *If the saturated fusion action system \mathfrak{X} on S is realized by G , $P \leq S$ is fully X -normalized if and only if $N_S(P; X) \in \text{Syl}_p(N_G(P; X))$.*

Similarly, P is fully X -centralized if and only if $Z_S(P; X) \in \text{Syl}_p(Z_G(P; X))$.

Proof. We prove the result for X -normalizers and note that the same argument works for X -centralizers.

If $N_S(P; X) = S \cap N_G(P; X)$ is Sylow in $N_G(P; X)$, then $|N_S(P; X)| \geq |N_S({}^gP; X)|$ for all $g \in G$ such that ${}^gP \leq S$.

Suppose now that $|N_S(P; X)| \geq |N_S({}^gP; X)|$ for all $g \in G$ such that ${}^gP \leq S$. Pick $T \in \text{Syl}_p(N_G(P; X))$, and $g \in G$ such that $T \cdot P \leq {}^gS$. Then we have $P^g \leq S$, and the fact that $N_S(P; X) \in \text{Syl}_p(N_G(P; X))$ implies that $N_S(P^g; X) \in \text{Syl}_p(N_G(P^g; X))$ (this uses the fact that the core of the action is normal in G). The assumption that P is fully X -normalized now implies that the orders of the X -normalizers in S of P and $P^g = {}^{g^{-1}}P$ are equal, which in turn forces $N_S(P; X) \in \text{Syl}_p(N_G(P; X))$, as desired. \square

The final ingredient to define saturation is the notion of the extender of an isomorphism.

Definition 2.2.6. The *extender* of $(\varphi, \sigma) \in \mathfrak{X}(P, Q)_{\text{iso}}$ is the subgroup of $N_S(P)$ given by

$$N_{(\varphi, \sigma)} = \{n \in N_S(P) \mid (\varphi c_n \varphi^{-1}, \sigma \ell_n \sigma^{-1}) \in \text{Aut}_S(Q; X)\}.$$

Definition 2.2.7 (Saturation Axioms for Fusion Action Systems). The abstract action \mathfrak{X} is *saturated* if

1. For any $P \leq S$, the following implications hold:

$$\begin{array}{ccc} P \text{ fully normalized} & \implies & P \text{ fully } X\text{-normalized} \\ \Downarrow & & \Downarrow \\ P \text{ fully centralized} & \implies & P \text{ fully } X\text{-centralized} \end{array}$$

2. If P is fully normalized, then

- $\text{Aut}_S(P) \in \text{Syl}_p(\mathcal{F}^{\mathfrak{X}}(P))$.
- $\text{Aut}_S(P; X) \in \text{Syl}_p(\mathfrak{X}(P))$.
- $\Sigma_{\mathfrak{X}}^S(P) \in \text{Syl}_p(\Sigma_{\mathfrak{X}}(P))$.

3. If P is fully X -normalized, then $\mathcal{F}_S(P)_0 \in \text{Syl}_p(\mathcal{F}^{\mathfrak{X}}(P)_0)$.

4. If P is fully centralized, then $\Sigma_{\mathfrak{X}}^S(P)_0 \in \text{Syl}_p(\Sigma_{\mathfrak{X}}(P)_0)$.

5. If Q is fully X -centralized and $(\varphi, \sigma) \in \mathfrak{X}(P, Q)_{\text{iso}}$, then there is some $(\tilde{\varphi}, \sigma) \in \mathfrak{X}(N_{(\varphi, \sigma)}, S)$ that extends (φ, σ) , i.e., such that $\varphi = \tilde{\varphi}|_P$.

We shall refer to Points 2-4 collectively as the *Sylow Axioms* and Point 5 as the *Extension Axiom* for fusion action systems.

Notation 2.2.8. Given a fusion action system \mathfrak{X} , $P \leq S$, and $Q \leq N_S(P)$, we shall denote by \overline{Q} the image of Q in $\mathfrak{X}(P)$; that is, $\overline{Q} = \{(c_q|_P, \ell_q) \mid q \in Q\}$.

It is obvious that the Saturation Axioms for fusion action systems are very similar to those of fusion systems. Let us note a further connection between the two.

Proposition 2.2.9. *If the abstract fusion action system \mathfrak{X} is saturated, the underlying fusion system $\mathcal{F}^{\mathfrak{X}}$ is as well.*

Proof. All of the saturation conditions are clear except for the Extension Axiom. Pick $\varphi \in \mathcal{F}^{\mathfrak{X}}(P, Q)_{\text{iso}}$ with Q fully centralized in $\mathcal{F}^{\mathfrak{X}}$. Recall that

$$N_{\varphi} = \{n \in N_S(P) \mid \varphi c_n \varphi^{-1} \in \text{Aut}_S(Q)\}.$$

If $(\varphi, \sigma) \in \mathfrak{X}(P, Q)$ is a morphism that lies over φ , we have

$$N_{(\varphi, \sigma)} = \{n \in N_S(P) \mid (\varphi c_n \varphi^{-1}, \sigma \ell_n \sigma^{-1}) \in \text{Aut}_S(Q; X)\}.$$

Clearly $N_{(\varphi, \sigma)} \leq N_{\varphi}$ for all possible choices of σ . If we can show that we have equality for some particular choice of σ , the Extension Axiom for the fusion action system \mathfrak{X} will imply the desired Axiom for the fusion system $\mathcal{F}^{\mathfrak{X}}$.

We can restate the problem as follows: Fix some σ such that $(\varphi, \sigma) \in \mathfrak{X}(P, Q)$, and let $N \leq \mathfrak{X}(Q)$ denote the group $(\varphi, \sigma) \circ \overline{N_{\varphi}} \circ (\varphi, \sigma)^{-1}$. Then N is a p -subgroup of $\mathfrak{X}(Q)$. If $N \leq \text{Aut}_S(Q; X)$, by definition $N_{\varphi} = N_{(\varphi, \sigma)}$ and we are done. The goal then becomes to show that there is some $\tau \in \Sigma_{\mathfrak{X}}(Q)_0$ such that ${}^{\tau}N \leq \text{Aut}_S(Q; X)$, for this will imply that $N_{\varphi} = N_{(\varphi, \tau\sigma)}$ and complete the proof. (Here and elsewhere we

shall ignore the distinction between an element $\sigma \in \Sigma_S$ and the corresponding morphism (id_Q, σ) that lies in \mathfrak{X} .)

For $n \in N_\varphi$, we have $\varphi c_n \varphi^{-1} \in \text{Aut}_S(Q)$, so $(\varphi c_n \varphi^{-1}, \sigma \ell_n \sigma^{-1})$ differs from an element of $\text{Aut}_S(Q; X)$ by an element of $\Sigma_{\mathfrak{X}}(Q)_0$. The claim is that $\text{Aut}_S(Q; X)$ is Sylow in $\Sigma_{\mathfrak{X}}(Q)_0 \cdot \text{Aut}_S(Q; X)$. We compute

$$\begin{aligned} [\text{Aut}_S(Q; X) : \Sigma_{\mathfrak{X}}(Q)_0 \cdot \text{Aut}_S(Q; X)] &= \frac{|\Sigma_{\mathfrak{X}}(Q)_0| |\text{Aut}_S(Q; X)|}{|\Sigma_{\mathfrak{X}}(Q)_0 \cap \text{Aut}_S(Q; X)| |\text{Aut}_S(Q; X)|} \\ &= [\Sigma_{\mathfrak{X}}(Q)_0 \cap \text{Aut}_S(Q; X) : \Sigma_{\mathfrak{X}}(Q)_0] \end{aligned}$$

By Axiom (4) for saturated fusion action systems, $\Sigma_{\mathfrak{X}}(Q)_0 \cap \text{Aut}_S(Q; X) = \Sigma_{\mathfrak{X}}^S(Q)_0$ is Sylow in $\Sigma_{\mathfrak{X}}(Q)_0$. Thus there is some $\tau \in \Sigma_{\mathfrak{X}}(Q)_0$ such that ${}^\tau N \leq \text{Aut}_S(Q; X)$, and the result is proved. \square

It is possible to define a morphism of fusion action systems. Basically, the obvious thing to try works, but as we shall make no use of the category of fusion action systems in this article, we shall not define morphisms here.

We conclude this section with some general properties of abstract fusion action systems, starting with a sort of converse to the Sylow Axioms of the Saturation Axioms for fusion action systems.

Proposition 2.2.10. *Fix a saturated fusion action system \mathfrak{X} and a subgroup $P \leq S$.*

(a) *The following are equivalent:*

1. *P is fully normalized in \mathfrak{X} .*
2. *P is fully centralized in \mathfrak{X} and $\text{Aut}_S(P) \in \text{Syl}_p(\mathcal{F}(P))$.*
3. *P is fully X -normalized in \mathfrak{X} and $\Sigma_{\mathfrak{X}}^S(P) \in \text{Syl}_p(\Sigma_{\mathfrak{X}}(P))$.*
4. *P is fully X -centralized in \mathfrak{X} and $\text{Aut}_S(P; X) \in \text{Syl}_p(\mathfrak{X}(P))$.*

(b) *P is fully centralized iff P is fully X -centralized and $\Sigma_{\mathfrak{X}}^S(P)_0 \in \text{Syl}_p(\Sigma_{\mathfrak{X}}(P)_0)$.*

(c) *P is fully X -normalized iff P is fully X -centralized and $\mathcal{F}_S(P)_0 \in \text{Syl}_p(\mathcal{F}(P)_0)$.*

Proof. Half of these implications are part of the Saturation Axioms, so prove only the remaining ones.

(a) $P \leq S$ is fully normalized if:

- P is fully centralized and $\text{Aut}_S(P) \in \text{Syl}_p(\mathcal{F}(P))$. In this case, the assumptions on P together with the inclusion of short exact sequences

$$\begin{array}{ccccccc} 1 & \longrightarrow & \Sigma_{\mathfrak{X}}(P)_0 & \longrightarrow & \mathfrak{X}(P) & \longrightarrow & \mathcal{F}(P) \longrightarrow 1 \\ & & \uparrow \text{Sylow} & & \uparrow & & \uparrow \text{Sylow} \\ 1 & \longrightarrow & \Sigma_{\mathfrak{X}}^S(P)_0 & \longrightarrow & \text{Aut}_S(P; X) & \longrightarrow & \text{Aut}_S(P) \longrightarrow 1 \end{array}$$

force $\text{Aut}_S(P; X) \in \text{Syl}_p(\mathfrak{X}(P))$. Then if Q is fully normalized and \mathcal{F} -conjugate to P via $(\varphi, \sigma) \in \mathfrak{X}(Q, P)_{\text{iso}}$, we have that $(\varphi, \sigma) \circ \text{Aut}_S(Q; X) \circ (\varphi, \sigma)^{-1} \leq \mathfrak{X}(P)$ is an inclusion of a p -subgroup. Thus there is $(\psi, \tau) \in \mathfrak{X}(P)$ such that $(\psi\varphi, \tau\sigma) \circ \text{Aut}_S(Q; X) \circ (\psi\varphi, \tau\sigma)^{-1} \leq \text{Aut}_S(P; X)$.

Note that $\text{Aut}_S(Q; X)$ is the image in $\mathfrak{X}(Q)$ of $N_S(Q)$, and that P 's being fully centralized implies that it is fully X -centralized. The Extension Axiom now gives the existence of $(\widetilde{\psi}\varphi, \tau\sigma) \in \mathfrak{X}(N_S(Q), N_S(P))$, from which it follows that $|N_S(Q)| = |N_S(P)|$ and P is fully normalized as well.

- P is fully X -normalized and $\Sigma_{\mathfrak{X}}^S(P) \in \text{Syl}_p(\Sigma_{\mathfrak{X}}(P))$. The inclusion of short exact sequences

$$\begin{array}{ccccccc} 1 & \longrightarrow & \mathcal{F}(P)_0 & \longrightarrow & \mathfrak{X}(P) & \longrightarrow & \Sigma_{\mathfrak{X}}(P) \longrightarrow 1 \\ & & \uparrow \text{Sylow} & & \uparrow & & \uparrow \text{Sylow} \\ 1 & \longrightarrow & \mathcal{F}_S(P)_0 & \longrightarrow & \text{Aut}_S(P; X) & \longrightarrow & \Sigma_{\mathfrak{X}}^S(P) \longrightarrow 1 \end{array}$$

shows that $\text{Aut}_S(P; X) \in \text{Syl}_p(\mathfrak{X}(P))$. The rest of the proof is the same as the first point.

- P is fully X -centralized and $\text{Aut}_S(P; X) \in \text{Syl}_p(\mathfrak{X}(P))$. The argument is the same as the end of the previous two.
- (b) P is fully centralized if P is fully X -centralized and $\Sigma_{\mathfrak{X}}^S(P)_0 \in \text{Syl}_p(\Sigma_{\mathfrak{X}}(P)_0)$: Observe that $|Z_S(P)| = |Z_S(P; X)| |\Sigma_{\mathfrak{X}}^S(P)_0|$. If $P \cong_{\mathcal{F}\mathfrak{X}} Q$, then we have $\Sigma_{\mathfrak{X}}(P)_0 \cong \Sigma_{\mathfrak{X}}(Q)_0$, so it is easy to see that the order of (the p -group) $Z_S(P)$ is maximized precisely when $\Sigma_{\mathfrak{X}}^S(P)_0 \in \text{Syl}_p(\Sigma_{\mathfrak{X}}(P)_0)$ and the order of $Z_S(P; X)$ is maximized. The result follows.
- (c) P is fully X -normalized if P is fully X -centralized and $\mathcal{F}_S(P)_0 \in \text{Syl}_p(\mathcal{F}(P)_0)$: The same argument as the previous paragraph applies, with the observation that $|N_S(P)| = |N_S(P; X)| |\mathcal{F}_S(P)_0|$.

□

Let us close with a basic reality-check for our Axioms: If these Saturation Axioms are “right,” the least we would expect is for some sort of Alperin Fusion Theorem to hold.

Proposition 2.2.11. *If $(\varphi, \sigma) \in \mathfrak{X}(P, Q)$ is an isomorphism such that Q is fully normalized, then there are morphisms $(\tilde{\varphi}, \sigma') \in \mathfrak{X}(N_S(P), S)$ and $(\psi, \tau) \in \mathfrak{X}(P)$ such that $\tilde{\varphi}|_P = \varphi \circ \psi$ and $\sigma' = \sigma \circ \tau$.*

Proof. By the assumption that Q is fully normalized, the Saturation Axioms state that $\text{Aut}_S(Q; X) \in \text{Syl}_p(\mathfrak{X}(Q))$, so the p -group $(\varphi, \sigma) \circ \text{Aut}_S(P; X) \circ (\varphi^{-1}, \sigma^{-1}) \leq \mathfrak{X}(Q)$ is subconjugate to $\text{Aut}_S(Q; X)$. Say

$$(\chi, v) \circ (\varphi, \sigma) \circ \text{Aut}_S(P; X) \circ (\varphi, \sigma)^{-1} \circ (\chi, v)^{-1} \leq \text{Aut}_S(Q; X)$$

and set $(\psi, \tau) = (\varphi^{-1}\chi\varphi, \sigma^{-1}v\sigma) \in \mathfrak{X}(P)$, so that

$$(\varphi, \sigma) \circ (\psi, \tau) \circ \text{Aut}_S(P; X) \circ (\psi, \tau)^{-1} \circ (\varphi, \sigma)^{-1} \leq \text{Aut}_S(Q; X).$$

Now the fact that Q is fully normalized and thus fully X -centralized implies that $(\varphi\psi, \sigma\tau) \in \mathfrak{X}(P, Q)$ has an extension $(\tilde{\varphi}, \sigma') \in \mathfrak{X}(N_{(\varphi\chi, \sigma\tau)}(P), S)$ by the Extension Axiom. Recall that

$$N_{(\varphi\chi, \sigma\tau)}(P) = \{s \in N_S(P) \mid (\varphi\chi \circ c_s \circ \chi^{-1}\varphi^{-1}, \sigma\tau \circ \ell_s \circ \tau^{-1}\sigma^{-1}) \in \text{Aut}_S(Q; X)\}$$

so that by construction $N_S(P) \leq N_{(\varphi\chi, \sigma\tau)}(P)$. The result follows. □

One could also prove analogous statements in the cases that the target is fully centralized or X -normalized, but we will not have need for such results.

Theorem 2.2.12 (Alperin Fusion Theorem for Fusion Action Systems). *If the fusion action system \mathfrak{X} is saturated, every morphism of \mathfrak{X} can be written as a composite of automorphisms of subgroups of S and inclusions.*

Proof. The proof goes by downward induction on the order of the source P . If $P = S$, then $(\varphi, \sigma) \in \mathfrak{X}(S)$, and there’s nothing to prove.

Therefore suppose that $P \leq S$. Without loss of generality, we may assume that $Q = \varphi(P)$ is fully normalized: Otherwise, pick Q' in the \mathcal{F} -conjugacy class of P that is fully normalized and an isomorphism $(\psi, \tau) \in \mathfrak{X}(P, Q')$. If the result is true for (ψ, τ) and $(\varphi, \sigma) \circ (\psi, \tau)^{-1}$, it clearly is for (φ, σ) as well.

Proposition 2.2.11 shows that if the target of (φ, σ) is fully normalized, then (φ, σ) can be composed with an element of $\mathfrak{X}(P)$ so that the resulting morphism extends to $\mathfrak{X}(N_S(P), S)$. But $N_S(P) \geq P$ as P is a proper subgroup, and the inductive hypothesis gives the rest. □

2.3 Faithful fusion action systems

Note that a fusion system \mathcal{F} is actually a special case of the notion of fusion action system, where the set X being acted on is a one-point set. In particular, fusion systems could be thought of as fusion action systems where the S -action is trivial.

In this section we examine in closer detail a fusion action system \mathfrak{X} at the opposite extreme to \mathcal{F} : We assume that X is faithful as an S -set. We will see that in this situation, the underlying saturated fusion system \mathcal{F} is always realizable by a finite group G , and moreover that G acts on X in such a way that $\mathfrak{X} = \mathfrak{X}_G$.

Let $G = \mathfrak{X}(1)$ be the group of permutations of X that appear in some morphism of \mathfrak{X} . The map $S \rightarrow G : s \mapsto \ell_s$ is an injection if we assume that X is faithful as an S -set, so we will identify S with its image in G . The Saturation Axioms for fusion actions imply that $S \in \text{Syl}_p(G)$, so the fusion system \mathcal{F}_G is saturated.

The following proposition shows that exotic fusion systems cannot be the underlying fusion system of a S -faithful fusion action system.

Proposition 2.3.1. *If \mathfrak{X} is a saturated fusion action system and S acts faithfully on X , then $\mathcal{F}^{\mathfrak{X}} = \mathcal{F}_G$.*

Proof. First we show that $\mathcal{F}^{\mathfrak{X}} \subseteq \mathcal{F}_G$. It suffices to show that any $\varphi \in \mathcal{F}^{\mathfrak{X}}(P, S)$ is realized by conjugation by an element of G . Pick some $(\varphi, \sigma) \in \mathfrak{X}(P, S)$, so that φ and σ are intertwined, so $\sigma \ell_p \sigma^{-1} = \ell_{\varphi(p)} \in \Sigma_X$ for all $p \in P$. This is exactly to say that σ conjugates P via φ , proving the first inclusion.

For the reverse inclusion, suppose that for some $\sigma \in G$ and $P \leq S$ we have ${}^\sigma P \leq S$; we must show that $c_\sigma \in \mathcal{F}^{\mathfrak{X}}(P, S)$. By the Extension Axiom, the morphism $(\text{id}_1, \sigma) \in \mathfrak{X}(1)$ extends to some $(\varphi, \sigma) \in \mathfrak{X}(N_{(\text{id}_1, \sigma)}, S)$, where $N_{(\text{id}_1, \sigma)} = \{n \in S \mid \sigma \circ \ell_n \circ \sigma^{-1} \in \text{Aut}_S(X)\}$. In other words, for all $n \in N_{(\text{id}_1, \sigma)}$, there is some (necessarily unique by faithfulness of the X -action) $s \in S$ so that $\sigma \circ \ell_n \circ \sigma^{-1} = \ell_s$. Therefore $N_{(\text{id}_1, \sigma)}$ is the maximal subgroup of S that is conjugated into S by σ . Thus $P \leq N_{(\text{id}_1, \sigma)}$, and if we can show that the assignment $n \mapsto s$ is equal to φ , the result will follow, as φ is by assumption a morphism of $\mathcal{F}^{\mathfrak{X}}$.

This final assertion follows again from the fact that φ and σ are intertwined: $\sigma \ell_n \sigma^{-1} = \ell_{\varphi(n)}$ for $n \in N_{(\text{id}_1, \sigma)}$ forces $s = \varphi(n)$ in the above notation. The result is proved. \square

In this situation, there is a natural action of $G = \mathfrak{X}(1) \leq \Sigma_X$ on X , and the following is immediate:

Corollary 2.3.2. *The fusion action system \mathfrak{X} is realized by the G -action on X ; i.e., $\mathfrak{X} = \mathfrak{X}_G$.*

Proof. The key point is that if the S -action on X is faithful, then any morphism $(\varphi, \sigma) \in \mathfrak{X}(P, Q)$ is actually determined by σ alone. This follows from the assumption that φ and σ are intertwined, so that for all $p \in P$, we have $\ell_{\varphi(p)} = \sigma \ell_p \sigma^{-1}$ along with the identification $p \leftrightarrow \ell_p$. In particular, $\mathfrak{X} \subseteq \mathfrak{X}_G$, since $(\varphi, \sigma) \in \mathfrak{X}(P, Q)$ implies that ${}^\sigma P \leq Q$ and $c_\sigma = \varphi$.

On the other hand, for any $\sigma \in N_G(P, Q)$ we have $(c_\sigma, \sigma) \in \mathfrak{X}_G(P, Q)$, and the claim is that this is also a morphism of $\mathfrak{X}(P, Q)$. This follows from the Extension Axiom of saturated fusion actions as in Proposition 2.3.1 and the observation that ${}^\sigma P \leq S$, implies $P \leq N_{(\text{id}_1, \sigma)}$. Thus $\mathfrak{X}_G \subseteq \mathfrak{X}$. \square

2.4 Core subsystems

In this section let \mathfrak{X} be a saturated fusion action system. Recall that $C \leq S$ denotes the core of the S action on X .

Definition 2.4.1. The *core fusion system* associated to \mathfrak{X} is the fusion system \mathcal{C} on C with morphisms $\mathcal{C}(P, Q) = \{\varphi \in \text{Inj}(P, Q) \mid (\varphi, \text{id}_X) \in \mathfrak{X}(P, Q)\}$.

Remark 2.4.2. There is a natural functor $\mathfrak{X} \rightarrow \mathcal{B}\Sigma_X$ that sends inclusions to the identity; this is precisely what the various maps π_Σ piece together to give. The fusion system \mathcal{C} can be thought of the kernel fusion action system of this functor, once we note that, as the action of \mathcal{C} on X is trivial, there is no distinction between a fusion system and a fusion action system.

The first goal of this section is that \mathcal{C} is saturated as a fusion system.

Proposition 2.4.3. *\mathcal{C} satisfies the following properties:*

1. *\mathcal{C} is strongly closed in $\mathcal{F}^{\mathfrak{X}}$.*
2. *For all $P, Q \leq C$, $\varphi \in \mathcal{C}(P, Q)$ and $\psi \in \mathcal{F}^{\mathfrak{X}}(C)$, we have $\psi\varphi\psi^{-1} \in \mathcal{C}(\psi P, \psi Q)$.*
3. *For all $P, Q \leq C$ and $\chi \in \mathcal{F}^{\mathfrak{X}}(P, Q)$, there are $\psi \in \mathcal{F}^{\mathfrak{X}}(C)$ and $\varphi \in \mathcal{C}(\psi(P), Q)$ such that $\chi = \varphi \circ \psi|_P^{\psi(P)}$.*

Proof.

1. Cf. Lemma 2.2.3.
2. Pick some $(\psi, \tau) \in \mathfrak{X}(C)$ that lies over $\psi \in \mathcal{F}^{\mathfrak{X}}(C)$. Then $(\psi, \tau)(\varphi, \text{id}_X)(\psi, \tau)^{-1} = (\psi\varphi\psi^{-1}, \text{id}_X)$, and the result follows.
3. Let $(\chi, v) \in \mathfrak{X}(P, Q)$ lie over $\chi \in \mathcal{F}^{\mathfrak{X}}(P, Q)$, and consider $(\text{id}_1, v) \in \mathfrak{X}(1)$. By the Extension Axiom for saturated fusion action systems, (id_1, v) extends to a map $(\tilde{\psi}, v) \in \mathfrak{X}(N_{(\text{id}_1, v)}, S)$, and it follows easily from the definition that $C \leq N_{(\text{id}_1, v)}$. Let $(\psi, v) \in \mathfrak{X}(C)$ be the restriction to C (which is necessarily an automorphism of C), and set $(\varphi, \text{id}_X) = (\chi, v) \circ (\psi, v)^{-1}|_{\psi(P)}^P \in \mathfrak{X}(\psi(P), Q)$. Projection to $\mathcal{F}^{\mathfrak{X}}$ shows now that ψ and φ have the desired property. □

The content of Proposition 2.4.3 is that \mathcal{C} is almost a normal subsystem of \mathcal{F} in the terminology of Puig, which we shall call an *invariant subsystem*, following Aschbacher. We say “almost” because in the definition of invariance, we would also need that \mathcal{C} is saturated as a fusion system, which is our current goal. We need the third point of Proposition 2.4.3 to prove saturation, but luckily the logic is not circular.

Corollary 2.4.4. *Every $P \leq C$ is fully normalized (resp. centralized) in \mathcal{C} if and only if P is fully X -normalized (resp. X -centralized) in \mathfrak{X} .*

Proof. Note that $N_S(P; X) = N_C(P)$ and $Z_S(P; X) = Z_C(P)$ by definition, which makes the “if” implications obvious. We prove the “only if” implication for normalizers; the same argument works for centralizers.

Suppose that P is fully normalized in \mathcal{C} , and let $Q \leq C$ be fully X -normalized and $\mathcal{F}^{\mathfrak{X}}$ -conjugate to P (such a Q must exist as C is strongly closed in $\mathcal{F}^{\mathfrak{X}}$). For any isomorphism $\chi \in \mathcal{F}^{\mathfrak{X}}(P, Q)_{\text{iso}}$, we can find $\psi \in \mathcal{F}^{\mathfrak{X}}(C)$ and $\varphi \in \mathcal{C}(\psi(P), Q)$ such that $\chi = \varphi \circ \psi|_P^{\psi(P)}$ by the third point of Proposition 2.4.3. As ψ extends to an automorphism of C , it sends the X -normalizer of P to the X -normalizer of $\psi(P)$. As φ is necessarily an isomorphism in \mathcal{C} , the assumption that $|N_C(P)| \geq |N_C(\psi(Q))|$ forces that $|N_S(P; X)| \geq |N_S(Q; X)|$. The assumption that Q is fully X -normalized forces equality, and the result is proved. □

Theorem 2.4.5. *\mathcal{C} is a saturated fusion system.*

Proof. Using the Saturation Axioms of [OS], it suffices to show that $\text{Aut}_C(C) \in \text{Syl}_p(\mathcal{C}(C))$ and that fully normalized subgroups satisfy the Extension Axiom.

To see that $\text{Aut}_C(C) \in \text{Syl}_p(\mathcal{C}(C))$: We’ve already noted that C is strongly closed in $\mathcal{F}^{\mathfrak{X}}$, and therefore is fully X -normalized. By the Saturation Axioms for fusion action systems, we have $\mathcal{F}_S^{\mathfrak{X}}(C)_0 \in \text{Syl}_p(\mathcal{F}^{\mathfrak{X}}(C)_0)$. It is easy to see from the definitions that $\mathcal{F}^{\mathfrak{X}}(C)_0 = \mathcal{C}(C)$, and we can describe $\mathcal{F}_S^{\mathfrak{X}}(C)_0$ as the subgroup of $\text{Aut}_S(C)$ consisting of those automorphisms that are induced by an element of S that acts trivially on X , so that $\mathcal{F}_S^{\mathfrak{X}}(C)_0 = \text{Aut}_C(C)$.

To check the Extension Axiom: If $P \leq C$ is fully normalized in \mathcal{C} , by Corollary 2.4.4 we see that P is fully X -normalized in $\mathcal{F}^{\mathfrak{X}}$, and thus fully X -centralized. Given an isomorphism $\varphi \in \mathcal{C}(Q, P)_{\text{iso}}$, we have $(\varphi, \text{id}_X) \in \mathfrak{X}(Q, P)$, and the extension condition for fusion actions gives a morphism $(\tilde{\varphi}, \text{id}_X) \in \mathfrak{X}(N_{(\varphi, \text{id}_X)}^{\mathfrak{X}}, C)$ where the target may be assumed to be C instead of S because C is strongly closed in $\mathcal{F}^{\mathfrak{X}}$. We wish to show

that $N_\varphi^{\mathcal{C}} = \{n \in N_C(Q) \mid \varphi c_n \varphi^{-1} \in \text{Aut}_K(P)\}$ is contained in $N_{(\varphi, \text{id}_X)}^{\mathfrak{X}}$, as then $\tilde{\varphi}|_{N_\varphi^{\mathcal{C}}}$ will give the desired extension of φ . We have

$$N_{(\varphi, \text{id}_X)}^{\mathfrak{X}} = \{n \in N_S(Q) \mid (\varphi \circ c_n \circ \varphi^{-1}, \ell_n) \in \text{Aut}_S(P; X)\}.$$

Since C acts trivially on S , for $n \in N_C(Q)$ such that $\varphi \circ c_n \circ \varphi^{-1} = c_k \in \text{Aut}_C(P)$ we have

$$(\varphi \circ c_n \circ \varphi^{-1}, \ell_n) = (c_k, \text{id}_X) = (c_k, \ell_k) \in \text{Aut}_C(P; X) \leq \text{Aut}_S(P; X)$$

and the result is proved. \square

Corollary 2.4.6. \mathcal{C} is a normal subsystem of $\mathcal{F}^{\mathfrak{X}}$ in the sense of [Asc].

Proof. Recall that in Aschbacher's definition of normality, in addition to the results of Proposition 2.4.3 and Theorem 2.4.5, we must also verify the technical condition that for all $\varphi \in \mathcal{C}(C)$ there is an extension $\tilde{\varphi} \in \mathcal{F}^{\mathfrak{X}}(C \cdot Z_S(C))$ such that $[\tilde{\varphi}, Z_S(C)] \leq Z(C)$, or equivalently $[\tilde{\varphi}, Z_S(C)] \leq C$.

If $\varphi \in \mathcal{C}(C)$ we have by definition $(\varphi, \text{id}_X) \in \mathfrak{X}(C)$. Since C is strongly closed in $\mathcal{F}^{\mathfrak{X}}$, it is in particular fully X -centralized, and so (φ, id_X) extends to some morphism in $\mathfrak{X}(N_{(\varphi, \text{id}_X)}, S)$. We claim that $Z_S(C) \leq N_{(\varphi, \text{id}_X)}$. Indeed, for $z \in Z_S(C)$, we have $(\varphi, \text{id}_X) \circ (c_z, \ell_z) \circ (\varphi, \text{id}_X)^{-1} = (\text{id}_C, \ell_z) = (c_z, \ell_z)$ since z acts trivially on C . Let $(\tilde{\varphi}, \text{id}_X)$ be the extension of (φ, id_X) to $C \cdot Z_S(C)$. As the strongly $\mathcal{F}^{\mathfrak{X}}$ -closed C is in the domain of $\tilde{\varphi}$, we in fact have $(\varphi, \text{id}_X) \in \mathfrak{X}(C \cdot Z_S(C))$.

Now we just need to show that $[\tilde{\varphi}, Z_S(C)] \leq C$: For all $z \in Z_S(C)$ we want $\tilde{\varphi}(z)z^{-1} \leq C$, or equivalently $\ell_{\tilde{\varphi}(z)}\ell_z^{-1} = \text{id}_X$. But since the pair $(\tilde{\varphi}, \text{id}_X)$ is intertwined, we have $\ell_{\tilde{\varphi}(z)} = \ell_z$, and the result follows. \square

We close this section by describing how the fusion action system \mathfrak{X} can be thought of as an extension of \mathcal{C} by the finite group $G = \mathfrak{X}(1)$. First note that we have the functors $\mathcal{C} \xrightarrow{\iota} \mathfrak{X} \xrightarrow{\pi_{\mathcal{T}}} \mathcal{T}_G$, where ι is “injective” in the sense that \mathcal{C} naturally sits as a subcategory of \mathfrak{X} , and $\pi_{\mathcal{T}}$ is “surjective” in the sense that every morphism of \mathcal{T}_G is in the image of $\pi_{\mathcal{T}}$. Just as extensions of the finite group N by H determines a morphism $H \rightarrow \text{Out}(N)$, would like to say that a fusion action system gives rise to an outer action of the “quotient” group G on the core subsystem \mathcal{C} . First we must understand what is meant by an “outer action.” We begin by recalling the definition of morphism of fusion system.

Definition 2.4.7. For fusion systems (S, \mathcal{F}) and (S', \mathcal{F}') , a group map $\alpha : S \rightarrow S'$ is *fusion-preserving* if there exists a functor $F_\alpha : \mathcal{F} \rightarrow \mathcal{F}'$ such that $F_\alpha(P) = \alpha P$ for all $P \leq S$ and the following diagram of groups commutes for all $\varphi \in \mathcal{F}(P, Q)$:

$$\begin{array}{ccc} P & \xrightarrow{\alpha} & \alpha P \\ \varphi \downarrow & & \downarrow F_\alpha(\varphi) \\ Q & \xrightarrow{\alpha} & \alpha Q \end{array}$$

Such a functor F_α is necessarily unique if it exists.

A *morphism of fusion systems* $(S, \mathcal{F}) \rightarrow (S', \mathcal{F}')$ is a fusion-preserving morphism $\alpha : S \rightarrow S'$. The set of all such morphisms is denoted $\text{Hom}(\mathcal{F}, \mathcal{F}')$, and in the case that $(S, \mathcal{F}) = (S', \mathcal{F}')$ the resulting group of fusion preserving automorphisms of S is written $\text{Aut}(\mathcal{F})$.

Remark 2.4.8. Implicit in Definition 2.4.7 is the easy result that if $\alpha \in \text{Hom}(\mathcal{F}, \mathcal{F}')$, $\alpha' \in \text{Hom}(\mathcal{F}', \mathcal{F}'')$, and $F_\alpha, F_{\alpha'}$ are the corresponding functors, then $F_{\alpha'\alpha} = F_{\alpha'}F_\alpha$ is the functor that shows that $\alpha'\alpha \in \text{Hom}(\mathcal{F}, \mathcal{F}'')$.

There is an easy way to check which automorphisms of the group S actually lie in $\text{Aut}(\mathcal{F})$. To express this, let us recall a piece of terminology from Remark 2.1.10:

Definition 2.4.9. Given a p -group S , subgroups $P, Q \leq R \leq S$, and an injective map $\gamma : R \rightarrow S$, the *translation along γ from P to Q* is the map

$$\begin{array}{ccc} t_\gamma|_P^Q : & \text{Hom}(P, Q) & \longrightarrow \text{Hom}(\gamma P, \gamma Q) \\ & \eta \longmapsto & \gamma \eta \gamma^{-1} \end{array}$$

In cases where there will be no confusion, we simply write t_γ for $t_\gamma|_P^Q$.

Proposition 2.4.10. *Given a fusion system \mathcal{F} on the p -group S and $\alpha \in \text{Aut}(S)$, we have $\alpha \in \text{Aut}(\mathcal{F})$ if and only if $t_\alpha(\mathcal{F}(P, Q)) = \mathcal{F}(\alpha P, \alpha Q)$ for all $P, Q \leq S$.*

Proof. If there is a functor F_α making (α, F_α) a morphism of \mathcal{F} , for all $\varphi \in \mathcal{F}(P, Q)$ we must have $F_\alpha(\varphi) = \alpha\varphi\alpha^{-1} = t_\alpha(\varphi) \in \mathcal{F}(\alpha P, \alpha Q)$. Thus if α extends to a functor of \mathcal{F} , it follows that $t_\alpha(\mathcal{F}(P, Q)) \subseteq \mathcal{F}(\alpha P, \alpha Q)$ for all $P, Q \leq S$. The fact that α is an automorphism of S forces equality.

Conversely, if we have $t_\alpha(\mathcal{F}(P, Q)) = \mathcal{F}(\alpha P, \alpha Q)$, the assignment $\varphi \mapsto t_\alpha\varphi$ can easily be seen to give the action of the desired F_α on morphisms. \square

Example 2.4.11. $\mathcal{F}(S) \subseteq \text{Aut}(\mathcal{F})$. In other words, if $\varphi \in \text{Aut}(S)$ is a morphism in \mathcal{F} , φ is actually fusion preserving. This follows immediately from Proposition 2.4.10 together with the Divisibility Axiom of fusion systems, which implies that the restriction of a morphism in \mathcal{F} to a subgroup also lies in \mathcal{F} .

Moreover, it is easy to see that $\mathcal{F}(S) \trianglelefteq \text{Aut}(\mathcal{F})$: If $\alpha \in \text{Aut}(\mathcal{F})$ and $\varphi \in \mathcal{F}(S)$, $\alpha\varphi\alpha^{-1} = t_\alpha(\varphi) \in \mathcal{F}(S)$ again by Proposition 2.4.10.

Definition 2.4.12. The *inner automorphism group* of \mathcal{F} is $\text{Inn}(\mathcal{F}) := \mathcal{F}(S)$. The *outer automorphism group* of \mathcal{F} is the quotient $\text{Out}(\mathcal{F}) := \text{Aut}(\mathcal{F})/\text{Inn}(\mathcal{F})$.

Theorem 2.4.13. *The fusion action system \mathfrak{X} determines a unique homomorphism $\kappa : G \rightarrow \text{Out}(\mathcal{C})$.*

Proof. An easy application of the Extension Axiom for saturated fusion action systems implies that each $\sigma \in G$ appears in the second coordinate of some $(\varphi, \sigma) \in \mathfrak{X}(C)$. As \mathcal{C} is normal in $\mathcal{F}^\mathfrak{X}$, $\varphi \in \text{Aut}(\mathcal{C})$. The indeterminacy of the assignment $\sigma \mapsto \varphi$ is measured by $\mathcal{C}(C)$, so the assignment $\sigma \mapsto [\varphi] \in \text{Out}(\mathcal{C})$ is the desired map. \square

2.5 Simplification of the Saturation Axioms

The Saturation Axioms of Definition 2.2.7 took their form from an analysis of the situation of a finite group acting on a finite set, just as the first set of Saturation Axioms for fusion systems in [BLO2] could be derived. The problem with the original list of Axioms is that it is both redundant and annoying to check in specific cases. The goal of this Subsection is to give a much shorter list of Saturation Axioms, which will be very similar in form to the simplified Axioms for saturated fusion systems introduced in [RS]. The form of the proof follows an unpublished manuscript by Bob Oliver.

Let \mathfrak{X} be a fusion action system associated to a given S action on X .

Definition 2.5.1. A subgroup $P \leq S$ is *fully automized* in \mathfrak{X} if $\text{Aut}_S(P; X) \in \text{Syl}_p(\mathfrak{X}(P))$. P is *receiving* if for all $(\varphi, \sigma) \in \mathfrak{X}(Q, P)_{\text{iso}}$ there exists an extension $(\tilde{\varphi}, \sigma) \in \mathfrak{X}(N_{(\varphi, \sigma)}, S)$.

Theorem 2.5.2. *The fusion action system \mathfrak{X} is saturated if and only if every \mathfrak{X} -conjugacy class of subgroups contains a fully automized, receiving subgroup.*

The proof of Theorem 2.5.2 will take the form of most of the lemmas in this Subsection. If \mathfrak{X} is saturated in the original definition then every subgroup that is fully normalized in \mathfrak{X} is both fully automized and receiving. We must therefore just show the other implication.

Lemma 2.5.3. *If $P \leq S$ is fully automized then $\Sigma_\mathfrak{X}^S(P)_0 \in \text{Syl}_p(\Sigma_\mathfrak{X}(P)_0)$ and $\mathcal{F}_S^\mathfrak{X}(P)_0 \in \text{Syl}_p(\mathcal{F}^\mathfrak{X}(P)_0)$.*

Proof. Consider the inclusion of exact sequences

$$\begin{array}{ccccccc} 1 & \longrightarrow & \Sigma_\mathfrak{X}(P)_0 & \longrightarrow & \mathfrak{X}(P) & \longrightarrow & \mathcal{F}^\mathfrak{X}(P) \longrightarrow 1 \\ & & \uparrow & & \uparrow & & \uparrow \\ 1 & \longrightarrow & \Sigma_\mathfrak{X}^S(P)_0 & \longrightarrow & \text{Aut}_S(P; X) & \longrightarrow & \mathcal{F}_S^\mathfrak{X}(P) \longrightarrow 1 \end{array}$$

The middle inclusion is Sylow if and only if the outer two are, so we have the first implication. The second follows from considering the exact sequences of the form $1 \longrightarrow \mathcal{F}^{\mathfrak{X}}(P)_0 \longrightarrow \mathfrak{X}(P) \longrightarrow \Sigma_{\mathfrak{X}}(P) \longrightarrow 1$. \square

Lemma 2.5.4. *If P is fully automized and receiving then P is fully normalized.*

Proof. Pick $Q \leq S$ fully normalized and some isomorphism $(\varphi, \sigma) \in \mathfrak{X}(Q, P)_{\text{iso}}$. We wish to show that, perhaps after modifying the choice of isomorphism (φ, σ) , we can extend this isomorphism to the whole of the normalizer of Q ; as any such map will send $N_S(Q)$ into $N_S(P)$, the choice of Q to be fully normalized will complete the proof.

Let $\overline{N_S(Q)}$ denote the image of $N_S(Q)$ in $\mathfrak{X}(Q)$. Then ${}^{(\varphi, \sigma)}\overline{N_S(Q)} \leq \mathfrak{X}(P)$ is a p -subgroup. As P is fully automized, there is some $(\psi, \tau) \in \mathfrak{X}(P)$ such that ${}^{(\psi, \tau)}\left({}^{(\varphi, \sigma)}\overline{N_S(Q)}\right) \leq \text{Aut}_S(P; X)$, so $N_S(Q) = N_{(\psi\varphi, \tau\sigma)}$. The fact that P is receiving implies the existence of an extension of $(\psi\varphi, \tau\sigma)$ to a morphism in $\mathfrak{X}(N_S(Q), N_S(P))_{\text{iso}}$, and the result is proved. \square

Lemma 2.5.5. *Suppose that P is receiving.*

1. *If $\Sigma_{\mathfrak{X}}^S(P)_0 \in \text{Syl}_p(\Sigma_{\mathfrak{X}}(P)_0)$, then P is fully centralized in \mathfrak{X} .*
2. *If $\mathcal{F}_S^{\mathfrak{X}}(P)_0 \in \text{Syl}_p(\mathcal{F}^{\mathfrak{X}}(P)_0)$, then P is fully X -normalized in \mathfrak{X} .*

Proof. The same proof works in both cases, so we shall just present the first. Pick $Q \leq S$ fully centralized and $(\varphi, \sigma) \in \mathfrak{X}(Q, P)_{\text{iso}}$. Then ${}^{(\varphi, \sigma)}\overline{Z_S(Q)}$ is a p -subgroup of $\Sigma_{\mathfrak{X}}(P)_0$. The condition on P forces there is some $(\text{id}_P, \tau) \in \Sigma_{\mathfrak{X}}(P)_0$ such that ${}^{(\text{id}_P, \tau)}\left({}^{(\varphi, \sigma)}\overline{Z_S(Q)}\right) \leq \Sigma_{\mathfrak{X}}^S(P)_0 \leq \text{Aut}_S(P; X)$. Thus $Z_S(Q) \leq N_{(\varphi, \tau\sigma)}$, so the fact that P is receiving implies that there is a morphism of \mathfrak{X} that sends $Z_S(Q)$ into $Z_S(P)$. The choice of Q to be fully centralized implies that P is as well. \square

Lemma 2.5.6. *If P is receiving then P is fully X -centralized in \mathfrak{X} .*

Proof. For any isomorphism $(\varphi, \sigma) \in \mathfrak{X}(Q, P)_{\text{iso}}$, it follows from the definition that $Z_S(Q; X) \leq N_{(\varphi, \sigma)}$, so any such extension of (φ, σ) will send $Z_S(Q; X)$ into $Z_S(P; X)$. In particular, if Q is fully X -centralized we get the result. \square

Proposition 2.5.7. *Suppose that $Q \leq S$ is fully automized and receiving and we are given $(\varphi, \sigma) \in \mathfrak{X}(P, Q)_{\text{iso}}$.*

1. *If P is fully X -centralized then P is receiving.*
2. *If P is fully X -normalized then P is receiving and $\mathcal{F}_S^{\mathfrak{X}}(P)_0 \in \text{Syl}_p(\mathcal{F}^{\mathfrak{X}}(P)_0)$.*
3. *If P is fully centralized then P is receiving and $\Sigma_{\mathfrak{X}}^S(P)_0 \in \text{Syl}_p(\Sigma_{\mathfrak{X}}(P)_0)$.*
4. *If P is fully normalized then P is fully automized and receiving.*

Proof. By the proof of Lemma 2.5.4, we may choose (φ, σ) such that there is an extension $(\tilde{\varphi}, \sigma) \in \mathfrak{X}(N_S(P), N_S(Q))$.

1. Pick any subgroup $R \leq S$ and isomorphism $(\psi, \tau) \in \mathfrak{X}(R, P)_{\text{iso}}$. As Q is receiving, the morphism $(\varphi\psi, \sigma\tau) \in \mathfrak{X}(R, Q)$ has an extension $(\tilde{\varphi}\psi, \sigma\tau) \in \mathfrak{X}(N_{(\varphi\psi, \sigma\tau)}, S)$. We claim that $N_{(\psi, \tau)} \leq N_{(\varphi\psi, \sigma\tau)}$ and $\tilde{\varphi}\psi(N_{(\psi, \tau)}) \leq \tilde{\varphi}(N_S(P))$: If we can show these to be true, $(\tilde{\varphi}, \sigma)^{-1} \circ (\tilde{\varphi}\psi, \sigma\tau)|_{N_{(\psi, \tau)}}$ will define our desired extension of (ψ, τ) .

If $n \in N_{(\psi, \tau)}$, by definition we have $(\psi c_n \psi^{-1}, \tau \ell_n \tau^{-1}) = (c_s, \ell_s)$ for some $s \in N_S(P)$. Then $((\varphi\psi) c_n (\varphi\psi)^{-1}, (\sigma\tau) \ell_n (\sigma\tau)^{-1}) = (\varphi c_s \varphi^{-1}, \sigma \ell_s \sigma^{-1})$, and we want to claim this is equal to $(c_{s'}, \ell_{s'})$ for some $s' \in N_S(Q)$. This is true if and only if $s \in N_{(\varphi, \sigma)}$, but we have chosen (φ, σ) such that the extender is all of $N_S(P)$, and thus we have proved the first claim.

To prove the second claim, note that $(\tilde{\varphi}, \sigma)$ sends $Z_S(P; X)$ into $Z_S(Q; X)$, and as P is fully X -centralized, we conclude that this is an isomorphism and therefore $Z_S(Q; X) \leq \tilde{\varphi}(N_S(P))$. Now for $n \in N_{(\psi, \tau)}$, we again have $(\psi c_n \psi^{-1}, \tau \ell_n \tau^{-1}) = (c_s, \ell_s)$. Then we have

$$(c_{\tilde{\varphi}\psi(n)}, \ell_{\tilde{\varphi}\psi(n)}) = (\varphi c_s \varphi^{-1}, \sigma \ell_s \sigma^{-1}) = (c_{\tilde{\varphi}(s)}, \ell_{\tilde{\varphi}(s)})$$

which implies that $\tilde{\varphi}\psi(n)$ differs from the element $\tilde{\varphi}(s) \in \tilde{\varphi}(N_S(P))$ by an element of $Z_S(Q; X)$, which we've already noted lies in the image of $\tilde{\varphi}$. Therefore $\tilde{\varphi}\psi(N_{(\psi, \tau)}) \leq \tilde{\varphi}(N_S(P))$, completing the proof.

2. To see that P is receiving, it suffices to show that P is fully X -centralized. As $\tilde{\psi}$ takes $N_S(P)$ into $N_S(Q)$, it restricts to an injection of $N_S(P; X)$ into $N_S(Q; X)$ and the assumption that P is fully X -centralized forces this to be an isomorphism. As $Z_S(P; X) \leq N_S(P; X)$, and similarly for Q , the inverse to the isomorphisms of X -normalizers takes the X -centralizer of Q into the X -centralizer of P . But Q is receiving, and hence fully X -centralized, so the X -centralizers are isomorphic, and P is fully X -centralized as well.

To verify the Sylow condition, note that we've already seen that Q is fully X -normalized and X -centralized, so we have

$$|\mathcal{F}_S^{\mathfrak{X}}(P)_0| = |N_S(P; X)|/|Z_S(P; X)| = |N_S(Q; X)|/|Z_S(Q; X)| = |\mathcal{F}_S^{\mathfrak{X}}(Q)_0|$$

But we've also seen that $\mathcal{F}_S^{\mathfrak{X}}(Q)_0 \in \text{Syl}_p(\mathcal{F}^{\mathfrak{X}}(Q)_0)$, so the isomorphism $\mathcal{F}^{\mathfrak{X}}(Q) \cong \mathcal{F}^{\mathfrak{X}}(P)$ gives the result.

The proofs of Items 3 and 4 are the same as that of Item 2, replacing every mention of X -normalizers with centralizers and normalizers, respectively. \square

We finish this section with another a seeming further simplification of the Saturation Axioms, though of course it will turn out to be equivalent. This version is the analogue of the definition in, for example, [OS] (I first became aware of this simplification of the Saturation Axioms for fusion systems by reading an unpublished paper of Stancu).

Theorem 2.5.8. *The fusion action system \mathfrak{X} is saturated if and only if S is fully automized and every fully normalized subgroup of S is receiving.*

Proof. If \mathfrak{X} is saturated then the original Axioms imply these weakened conditions, so we must just prove the opposite implication. Using the simplified Axioms of Theorem 2.5.2, it suffices to show that every fully normalized subgroup is fully automized.

Suppose otherwise, and let $P \leq S$ be of maximal order with the property that P is fully normalized and $\text{Aut}_S(P; X) \notin \text{Syl}_p(\mathfrak{X}(P))$. By the assumption that S itself is fully automized, we must have that P is a proper subgroup. Then there is some p -power element $(\varphi, \sigma) \in \mathfrak{X}(P) - \text{Aut}_S(P; X)$ that normalizes $\text{Aut}_S(P; X)$. We claim that $N_{(\varphi, \sigma)} = N_S(P)$. Indeed, if $n \in N_S(P)$, we have $(\varphi c_n \varphi^{-1}, \sigma \ell_n \sigma^{-1}) = (c_{n'}, \ell_{n'})$ for some other $n' \in N_S(P)$ by since (φ, σ) normalizes $\text{Aut}_S(P; X)$, which proves the claim.

Now by assumption P is receiving, so there is some extension $(\tilde{\varphi}, \sigma) \in \mathfrak{X}(N_S(P))$ of (φ, σ) . We may assume without loss of generality that $(\tilde{\varphi}, \sigma)$ has p -power order, as (φ, σ) does. Pick $(\psi, \tau) \in \mathfrak{X}(N_S(P), Q)_{\text{iso}}$ such that Q is fully normalized in \mathfrak{X} . Then $^{(\psi, \tau)}(\tilde{\varphi}, \sigma) \in \mathfrak{X}(Q)$ is a p -element, and as Q has order strictly larger than P , we conclude that (by adjusting (ψ, τ) as necessary), $^{(\psi, \tau)}(\tilde{\varphi}, \sigma) = (c_n, \ell_n) \in \text{Aut}_S(Q; X)$ for some $n \in N_S(Q)$. Then $c_n \psi(P) = \psi \tilde{\varphi}(P) = \psi(P)$ since $\tilde{\varphi}(P) = \varphi(P) = P$, so $n \in N_S(\psi(P))$. Since $N_S(\psi(P)) \leq \psi(N_S(P))$ (because Q was chosen to be fully normalized), we can therefore consider the element $m = \psi^{-1}(n) \in N_S(P)$, and we claim that $(\tilde{\varphi}, \sigma) = (c_m, \ell_m) \in \mathfrak{X}(N_S(P))$. We calculate for $a \in N_S(P)$

$$mam^{-1} = \psi^{-1}(n\psi(a)n^{-1}) = \psi^{-1}(\psi \tilde{\varphi} \psi^{-1} \psi(a)) = \tilde{\varphi}(a)$$

The fact that (ψ, τ) is an intertwined pair implies that $\tau \ell_m \tau^{-1} = \ell_{\psi(m)} = \ell_n$. On the other hand, we've assumed that $\tau \sigma \tau^{-1} = \ell_n$, so the claim is verified.

But this is a contradiction, for we have realized (φ, σ) (after restriction to P) as a morphism induced by an element of S . We conclude that $\text{Aut}_S(P; X) \in \text{Syl}_p(\mathfrak{X}(P))$, and thus \mathfrak{X} is saturated. \square

It should be noted that the difference between the proof of Theorem 2.5.8 and the analogous result for fusion systems is almost purely formal; the above proof looks nearly identical to that found in [Lin]. The same phenomenon will reappear in the next Subsection, as well as later in the paper.

2.6 Stabilizer subsystems and K -normalizers

The motivation for this Subsection is to describe the “stabilizer” fusion action subsystem of a well-chosen point of X ; we will return to the question of why this is a worthwhile goal in Subsection 3.5. In particular, we shall show that these stabilizer subsystems are saturated, and thus the underlying fusion systems are as well. On the way to proving this fairly narrow result, we shall have to introduce an analogue of Puig’s notion of K -normalizers, and end up proving a much more general Theorem.

Definition 2.6.1. The fusion action system \mathfrak{X} is *transitive* if the natural group action of $\mathcal{G} = \mathfrak{X}(1)$ on X is transitive.

Definition 2.6.2. Given a transitive fusion action system \mathfrak{X} and a point $x \in X$, the *stabilizer fusion action system* is the fusion action system \mathfrak{X}_x on S_x whose action on X is given by restriction of the original S -action, and whose morphisms are given by

$$\mathfrak{X}_x(P, Q) = \{(\varphi, \sigma) \in \mathfrak{X}(P, Q) \mid \sigma(x) = x\}.$$

Definition 2.6.3. If \mathfrak{X} is a transitive fusion action system, then $x \in X$ is *fully stabilized* if $|S_x| \geq |S_y|$ for all $y \in X$.

Theorem 2.6.4. *If $x \in X$ is fully stabilized in the saturated fusion action system \mathfrak{X} , the stabilizer subsystem \mathfrak{X}_x is also saturated.*

Instead of trying to prove Theorem 2.6.4 directly, we shall instead generalize the result twice.

Definition 2.6.5. Given a fusion action system \mathfrak{X} and a subgroup $H \leq G = \mathfrak{X}(1)$, the *preimage fusion action system of H in \mathfrak{X}* is the fusion action system \mathfrak{X}_H on the subgroup $T = \ell^{-1}(S \cap H) \leq S$, with action on X the restriction of the S -action, and with morphisms given by

$$\mathfrak{X}_H(P, Q) \{(\varphi, \sigma) \in \mathfrak{X}(P, Q) \mid \sigma \in H\}$$

Theorem 2.6.6. *If \mathfrak{X} is a saturated fusion action system and $H \leq \mathfrak{X}(1)$ is chosen such that $\ell_T \in \text{Syl}_p(H)$, the preimage fusion action system \mathfrak{X}_H is saturated.*

If we denote by Σ_{X-x} the subgroup of Σ that fixes x , it is immediate from the definition that the preimage fusion action system $\mathfrak{X}_{\Sigma_{X-x}}$ is just the stabilizer subsystem \mathfrak{X}_x . Thus Theorem 2.6.4 will follow from Theorem 2.6.6 with the aid of the following lemma.

Lemma 2.6.7. *If the fusion action system \mathfrak{X} is transitive then $x \in X$ is fully stabilized if and only if $\ell_{S_x} \in \text{Syl}_p(\mathfrak{X}(1) \cap \Sigma_{X-x})$.*

Proof. As the core C is contained in all stabilizers S_y for $y \in X$ and $|\ell_{S_y}| = |S_y|/|C|$, it follows immediately that x is fully stabilized if and only if $|\ell_{S_x}|$ is maximal among the orders of the groups ℓ_{S_y} . Moreover, the Sylow Axioms for saturation imply that $\ell_S \in \text{Syl}_p(\mathfrak{X}(1))$, and clearly $\ell_{S_y} = \ell_S \cap (\mathfrak{X}(1)_y)$, so we find ourselves in the following situation:

Let G be a finite group that acts transitively on the finite set X . For $S \in \text{Syl}_p(G)$ and $x \in X$, we have $|S_x| \geq |S_y|$ for all $y \in X$ if and only if $S_x \in \text{Syl}_p(G_x)$. This result is easy to see, but we include the proof for the sake of completeness.

If $x \in X$ is such that the order of S_x is maximal, let T be a Sylow subgroup of G_x that contains S_x . As T is a p -subgroup of G and $S \in \text{Syl}_p(G)$, there is some $g \in G$ such that ${}^gT \leq S$. We have ${}^gG_x = G_{g \cdot x}$, so ${}^gT \leq S_{g \cdot x}$. The assumption on the maximal order of the stabilizer of x then implies that $|{}^gT| \leq |S_x|$, from which the assumption that $S_x \leq T$ implies that $S_x = T \in \text{Syl}_p(G_x)$.

Conversely, if $S_x \in \text{Syl}_p(G_x)$, for any $y \in X$, pick $g \in G$ such that $g \cdot y = x$. Then the fact that ${}^g G_y = G_x$ implies that ${}^g S_y$ is a p -subgroup of G_x , and hence subconjugate to S_x by the Sylow assumption. Thus $|S_y| \leq |S_x|$, as desired. \square

We digress briefly from the flow of this section to note a corollary that intuitively “should” be true.

Corollary 2.6.8. *Let \mathfrak{X} be a transitive saturated fusion action system and x a fully stabilized point of X . Then for every $y \in X$, the stabilizer S_y is $\mathcal{F}^\mathfrak{X}$ -subconjugate to S_x . In particular, the stabilizers of distinct fully stabilized points of X are isomorphic.*

Proof. First note that if (φ, σ) is a morphism of \mathfrak{X} such that $\sigma(y) = x$ and $s \in S_y$, then $\varphi(s)$ lies in S_x when defined. This is simply a restatement of the fact that (φ, σ) is an intertwined pair: $\varphi(s) \cdot x = \varphi(s) \cdot \sigma(y) = \sigma(s \cdot y) = \sigma(y) = x$.

Thus the result will follow from the Extension Axiom if we can find some $\sigma \in \mathfrak{X}(1)$ such that $\sigma(y) = x$ and $S_y \leq N_{(1, \sigma)}$. As \mathfrak{X} is assumed to be transitive, there exists a $\sigma \in \mathfrak{X}(1)$ such that $\sigma(y) = x$. Then the group $H := \{\sigma \circ \ell_s \circ \sigma^{-1} \mid s \in S_y\}$ is a p -subgroup of $\mathfrak{X}(1)_x$. By Proposition 3.5.5 (a), $\ell_{S_x} \in \text{Syl}_p(\mathfrak{X}(1)_x)$, so without loss of generality we may assume that we have chosen σ such that $H \leq \mathfrak{X}(1)_x$ and $\sigma(y) = x$. But then $S_y \leq N_{(1, \sigma)}$, as desired. \square

Returning to our overall goal for this section, we must generalize some more.

Definition 2.6.9. For X an S -set and $P \leq S$, let $\text{Aut}(P; X)$ denote the group of all pairs

$$\{(\varphi, \sigma) \mid \varphi \in \text{Aut}(P), \sigma \in \Sigma_X, \varphi \text{ and } \sigma \text{ are intertwined}\}.$$

Similarly, let $\text{Aut}_S(P; X)$ denote the subgroup consisting of all pairs of the form (c_s, ℓ_s) for $s \in N_S(P)$.

Definition 2.6.10. Let K be a subgroup of $\text{Aut}(P; X)$. The K -normalizer of P in S is the group

$$N_S^K(P) = \{n \in N_S(P) \mid (c_n, \ell_n) \in K\}$$

We denote by $\text{Aut}_S^K(P)$ the group $\text{Aut}_S(P; X) \cap K$, so that $\text{Aut}_S^K(P) \cong N_S^K(P)/Z_S(P; X)$. Similarly set $\text{Aut}_{\mathfrak{X}}^K(P) := \mathfrak{X}(P) \cap K$.

Definition 2.6.11. Given a fusion action system \mathfrak{X} , a subgroup $P \leq S$, and some $K \leq \text{Aut}(P; X)$, the K -normalizer subsystem of P in \mathfrak{X} is the fusion action subsystem $\mathfrak{N} := N_{\mathfrak{X}}^K(P)$ on $N_S^K(P)$ whose morphisms are given by

$$\mathfrak{N}(Q, R) = \{(\varphi, \sigma) \in \mathfrak{X}(P, Q) \mid \exists (\tilde{\varphi}, \sigma) \in \mathfrak{X}(PQ, PR) \text{ s.t. } \tilde{\varphi}|_Q = \varphi \text{ and } (\tilde{\varphi}|_P, \sigma) \in K\}$$

The goal is to relate this notion to preimage fusion action systems, and thus back to stabilizer subsystems. Of course, we must still find conditions that will make K -normalizer subsystems saturated.

Note that if (φ, σ) is any intertwined morphism pair from P to Q , and $K \leq \text{Aut}(P; X)$, then ${}^{(\varphi, \sigma)}K := (\varphi, \sigma) \circ K \circ (\varphi, \sigma)^{-1}$ is a subgroup of $\text{Aut}(Q; X)$. We may thus compare the respective normalizer groups $N_S^K(P)$ and $N_S^{({}^{(\varphi, \sigma)}K)}(Q)$, or at least their orders.

Definition 2.6.12. For a fusion action system \mathfrak{X} on S , $P \leq S$, and $K \leq \text{Aut}(P; X)$, we say that P is *fully K -normalized in \mathfrak{X}* if $|N_S^K(P)| \geq |N_S^{({}^{(\varphi, \sigma)}K)}(Q)|$ for all $(\varphi, \sigma) \in \mathfrak{X}(P, Q)_{\text{iso}}$.

Let us record a basic fact about K -normalizers to get a better sense of how their use will play out.

Lemma 2.6.13. *If $(\varphi, \sigma) \in \mathfrak{X}(P, Q)_{\text{iso}}$ has some extension $(\tilde{\varphi}, \sigma) \in \mathfrak{X}(\tilde{P}, S)$, then for any $n \in \tilde{P} \cap N_S^K(P)$ we have $\tilde{\varphi}(n) \in N_S^{({}^{(\varphi, \sigma)}K)}(Q)$.*

Proof. If $n \in N_S^K(P)$ then $(c_n, \ell_n) \in K$, so it will suffice to show that if $\tilde{\varphi}(n)$ is defined then $(c_{\tilde{\varphi}(n)}, \ell_{\tilde{\varphi}(n)}) = (\varphi, \sigma)(c_n, \ell_n)(\varphi, \sigma)^{-1}$ on Q . Any $q \in Q$ can be written uniquely as $\varphi(p)$ for $p \in P$, and the fact that $\tilde{\varphi}$ is an extension of φ implies that

$$c_{\tilde{\varphi}(n)}(q) = \tilde{\varphi}(n)\tilde{\varphi}(p)\tilde{\varphi}(n)^{-1} = \tilde{\varphi}(n p n^{-1}) = \varphi(n p n^{-1}) = (\varphi \circ c_n \circ \varphi^{-1})(q)$$

while the fact that $\ell_{\tilde{\varphi}(n)} = \sigma \circ \ell_n \circ \sigma^{-1}$ follows directly from the assumption that $(\tilde{\varphi}, \sigma)$ is an intertwined pair. \square

Corollary 2.6.14. *If $(\varphi, \sigma) \in \mathfrak{X}(P, Q)_{\text{iso}}$ extends to $(\tilde{\varphi}, \sigma) \in \mathfrak{X}(P \cdot N_S^K(P), S)$, then if P is fully K -normalized we have that Q is fully $^{(\varphi, \sigma)}K$ -normalized.*

Proof. Lemma 2.6.13 implies that in this situation $N_S^K(P)$ injects into $N_S^{(\varphi, \sigma)K}(Q)$; the assumption that P is fully K -normalized implies that this is an isomorphism and the result follows. \square

Theorem 2.6.15. *With the notation as above, if P is fully K -normalized in the saturated fusion action system \mathfrak{X} , the K -normalizer subsystem \mathfrak{N} is saturated as well.*

Before we begin the proof of Theorem 2.6.15 let us see how this implies Theorem 2.6.6. Given $H \leq \mathfrak{X}(1)$, we can view H as a subgroup of $\text{Aut}(1; X)$, and thus consider the H -normalizer $N_S^H(1)$. By definition, this is the subgroup of X whose elements n are characterized by $\ell_n \in H$. Therefore the H -normalizer fusion action subsystem $\mathfrak{N} = N_{\mathfrak{X}}^H(1)$ is a fusion action system on the same group T on which the preimage fusion action subsystem \mathfrak{X}_H is based. Moreover, $(\varphi, \sigma) \in \mathfrak{N}(P, Q)$ if and only if $\sigma \in H$, so we see that in fact these two subsystems are equal. We are left to understand in what circumstances the identity subgroup is fully H -normalized in \mathfrak{X} , which the following Lemma will demonstrate is exactly the condition placed on H in Theorem 2.6.6:

Lemma 2.6.16. *Given a saturated fusion action system \mathfrak{X} , $P \leq S$, and $K \leq \text{Aut}(P; X)$, then P is fully K -normalized in \mathfrak{X} if and only if P is fully X -centralized in $\mathcal{F}^{\mathfrak{X}}$ and $\text{Aut}_S^K(P; X) \in \text{Syl}_p(\text{Aut}_{\mathfrak{X}}^K(P; X))$.*

Proof. First suppose that P is fully X -centralized in $\mathcal{F}^{\mathfrak{X}}$ and $\text{Aut}_S^K(P; X) \in \text{Syl}_p(\text{Aut}_{\mathfrak{X}}^K(P; X))$. For any $(\varphi, \sigma) \in \mathfrak{X}(P, Q)_{\text{iso}}$ it is immediate that $\text{Aut}_{\mathfrak{X}}^K(P; X) \cong \text{Aut}_{\mathfrak{X}}^{(\varphi, \sigma)K}(Q; X)$. The second assumption on P then implies that $|\text{Aut}_S^K(P; X)| \geq |\text{Aut}_{\mathfrak{X}}^{(\varphi, \sigma)K}(Q; X)|$, and the assumption that P is fully X -centralized lets us write

$$|N_S^K(P)| = |Z_S(P; X)| \cdot |\text{Aut}_S^K(P; X)| \geq |Z_S(Q; X)| \cdot |\text{Aut}_{\mathfrak{X}}^{(\varphi, \sigma)K}(Q; X)| = |N_S^{(\varphi, \sigma)K}(Q; X)|$$

so that P is fully K -normalized.

Now suppose that P is fully K -normalized, and pick an $\mathcal{F}^{\mathfrak{X}}$ -conjugate subgroup Q that is fully normalized. Without loss of generality [cite result] we may assume that we have an isomorphism $(\varphi, \sigma) \in \mathfrak{X}(P, Q)_{\text{iso}}$ that extends to $(\tilde{\varphi}, \sigma) \in \mathfrak{X}(N_S(P), N_S(Q))$. Lemma 2.6.13 implies that $\tilde{\varphi}(N_S^K(P)) \leq N_S^{(\varphi, \sigma)K}(Q)$, so by the assumption that P is fully K -normalized we must have that $\tilde{\varphi}|_{N_S^K(P)}$ is actually an isomorphism. The X -centralizer of any subgroup is contained in any K -normalizers (indeed, $Z_S(P; X) = N_S^{\{\text{id}_P\}}(P)$ is the minimal K -normalizer group), we see that $\tilde{\varphi}$ induces an isomorphism on X -centralizers. The fact that Q is fully X -centralized implies then that P is as well.

To verify the Sylow condition, note that $(\varphi, \sigma) \circ \text{Aut}_S^K(Q; X) \circ (\varphi, \sigma)^{-1}$ is a p -subgroup of $\text{Aut}_{\mathfrak{X}}^{(\varphi, \sigma)K}(Q; X) \leq \mathfrak{X}(P)$. As Q was taken to be fully normalized, $\text{Aut}_S(Q; X) \in \text{Syl}_p(\mathfrak{X}(P))$, so we conclude that there is some $(\psi, \tau) \in \mathfrak{X}(Q)$ such that

$$(\varphi, \sigma) \circ \text{Aut}_S^K(Q; X) \circ (\varphi, \sigma)^{-1} \leq (\psi, \tau) \circ \text{Aut}_S(Q; X) \circ (\psi, \tau)^{-1} \cap \text{Aut}_{\mathfrak{X}}^{(\varphi, \sigma)K}(Q; X) \in \text{Syl}_p(\text{Aut}_{\mathfrak{X}}^{(\varphi, \sigma)K}(Q; X))$$

so that

$$\text{Aut}_S(Q; X) \cap \text{Aut}_{\mathfrak{X}}^{((\varphi, \sigma)K)^{(\psi, \tau)}}(Q; X) = \text{Aut}_S^{((\varphi, \sigma)K)^{(\psi, \tau)}}(Q; X) \in \text{Syl}_p\left(\text{Aut}_{\mathfrak{X}}^{((\varphi, \sigma)K)^{(\psi, \tau)}}(Q; X)\right)$$

Now the assumption that P is fully K -normalized implies that $|\text{Aut}_S^K(Q; X)| \geq \left| \text{Aut}_S^{((\varphi, \sigma)K)^{(\psi, \tau)}}(Q; X) \right|$, from which it follows that $\text{Aut}_S^K(Q; X) \in \text{Syl}_p(\text{Aut}_X^K(Q; X))$. \square

As the identity subgroup is always fully X -centralized, we see that the primage fusion action subsystem \mathfrak{X}_H is saturated precisely when the condition of Theorem 2.6.6 is met, i.e., when $\ell_T \in \text{Syl}_p(H)$.

Corollary 2.6.17. *If $(\varphi, \sigma) \in \mathfrak{X}(P, Q)_{\text{iso}}$ and $K \leq \text{Aut}(P; X)$ are given such that Q is fully $(\varphi, \sigma)K$ -normalized. If \mathfrak{X} is saturated, then there exist $(\psi, \tau) \in \mathfrak{X}(P \cdot N_S^K(P), S)$ and $(\chi, \nu) \in \mathfrak{X}(P)$ such that $(\psi|_P, \tau) = (\varphi\chi, \sigma\nu)$.*

Proof. Let $N = \overline{N_S^K(P)}$ denote the image of $N_S^K(P)$ in $\mathfrak{X}(P)$, so that $(\varphi, \sigma) \circ N \circ (\varphi, \sigma)^{-1}$ is a p -subgroup of $\text{Aut}_X^{(\varphi, \sigma)K}(Q; X)$. By Lemma 2.6.16, the assumption on Q implies that $\text{Aut}_S^{(\varphi, \sigma)K}(Q; X) \in \text{Syl}_p\left(\text{Aut}_X^{(\varphi, \sigma)K}(Q; X)\right)$, so there is some $(\chi', \nu') \in \text{Aut}_X^{(\varphi, \sigma)K}(Q; X)$ such that $(\chi'\varphi, \nu'\sigma) \circ N \circ (\chi'\varphi, \nu'\sigma)^{-1} \leq \text{Aut}_S^{(\varphi, \sigma)K}(Q; X)$. Thus $N \leq N_{(\chi'\varphi, \nu'\sigma)}$, so using the fact that Q is fully X -centralized by Lemma 2.6.16 and thus receiving, we have the existence of an extension $(\psi, \tau) \in \mathfrak{X}(P \cdot N_S^K(P), S)$ of $(\chi'\varphi, \nu'\sigma)$. Finally, define $(\chi, \nu) := (\varphi^{-1}\chi'\varphi, \sigma^{-1}\nu'\sigma)$ to get the final result. \square

We now turn to the proof of Theorem 2.6.15, by which we mean that one can now observe that the setup of K -normalizers of fusion action systems is so closely analogous to that of fusion systems that in fact Puig's original proof of the analogous result in [Pui1], or, for example, Linckelmann's version of the same result in [Lin] goes through mutatis mutandis (meaning basically appending a permutation of X to every morphism of the fusion system). Note that in order to appeal to Linckelmann's proof we make use of the simplified Saturations Axioms of Theorem 2.5.8. As reproducing the proofs almost word for word is not particularly enlightening in this context, the interested reader should refer to the original sources.

3 p -local finite group actions

3.1 The ambient case

Let us return for this Subsection to the situation where we are given a finite group G , a Sylow $S \in \text{Syl}_p(G)$, the fusion system \mathcal{F}_G , and X a G -set. We wish to reconstruct the homotopy type of the Borel construction $B_G X := EG \times_G X$, at least up to p -completion, with a minimum of p' -data. This section reproduces some results of [BLO1] in the context of fusion action systems arising from ambient groups.

The game we want to play in reconstructing the p -completed homotopy type of $B_G X$ is to look for a new category that will both allow us to construct $B_G X_p^\wedge$ but that does not contain too much extra information. In some sense, the transporter system $\mathcal{T}_G = \mathcal{T}_S(G)$ contains all the information of G that we care about (for instance, the natural functor $\mathcal{B}G \rightarrow \mathcal{T}_G$ that sends the unique object of $\mathcal{B}G$ to the subgroup $\{1\}$ induces a homotopy equivalence $|\mathcal{T}_G| \simeq BG$). The problem of course is that \mathcal{T}_G contains too much information. The goal then becomes figuring out the right amount of data of \mathcal{T}_G to forget and still be able to understand the Borel construction.

The first way of forgetting information of \mathcal{T}_G is to consider full subcategories whose sets of objects are closed under G -conjugacy and overgroups. In other words, we simply throw out all sufficiently small subgroups and the information of their Hom-sets. Exactly which subgroups we will allow must depend somehow on the fusion data of G and the action of G on X :

Definition 3.1.1. A p -subgroup $P \leq S$ is *p -centric at X* if $Z(P; X) \in \text{Syl}_p(Z_G(P; X))$. Equivalently, P is *p -centric at X* if $Z_G(P; X) = Z(P; X) \times Z'_G(P; X)$ for some (necessarily unique) p' -group $Z'_G(P; X)$.

We shall generally call such a subgroup *X -centric* and omit mentioning the prime p . This frees up our nomenclature so we can recall

Definition 3.1.2. A p -subgroup $P \leq S$ is p -centric if $Z(P) \in \text{Syl}_p(Z_G(P))$, or equivalently if $Z_G(P) = Z(P) \times Z'_G(P)$ for some (again, unique) p' -group $Z'_G(P)$. This is equivalent to saying that P is X -centric for $X = *$ the trivial S -set.

Note that the condition of being X -centric is determined purely by fusion data and makes no reference to the fusion action system \mathfrak{F}_G . This motivates the following definition:

Definition 3.1.3. Given a saturated fusion system \mathcal{F} on S and an S -set X , a subgroup $P \leq S$ is \mathcal{F} -centric at X if $Z(Q; X) = Z_S(Q : X)$ for all $Q \cong_{\mathcal{F}} P$.

Remark 3.1.4. Actually, if we're fixing the fusion system on S without going straight to a fusion action system, we really should require that the S -set X is \mathcal{F} -stable. By this we mean that for all \mathcal{F} -conjugate subgroups P and Q we have $|X^P| = |X^Q|$, or a couple of other equivalent conditions. If we give ourselves an ambient group G , or even a fusion action system, the \mathcal{F} -stability condition is automatically realized, so we shall not concern ourselves with it too much.

In the presence of an ambient group G , the two notions of X -centricity coincide:

Proposition 3.1.5. If G is a finite group that acts on X and $S \in \text{Syl}_p(G)$, then a subgroup $P \leq S$ is p -centric at X if and only if it is \mathcal{F}_G -centric at X .

Proof. First suppose that P is p -centric at X . As $Z_G(P; X) = Z(P; X) \times Z'_G(P; X)$ for some uniquely defined p' -subgroup $Z'_G(P; X)$, it is clear that $Z_S(P; X) = Z(P; X)$. If $g \in G$ is such that ${}^gP \leq S$, the fact that $Z_G({}^gP; X) = {}^gZ_G(P; X)$ immediately shows that $Z_G({}^gP; X) = Z({}^gP; X) \times {}^gZ'_G(P; X)$, and we have the desired conclusion for gP .

Conversely, suppose that P is \mathcal{F}_G -centric at X , and pick $T \in \text{Syl}_p(Z_G(P; X))$ such that $Z(P; X) \leq T$. As $S \in \text{Syl}_p(G)$, there is some $g \in G$ such that ${}^gP \leq {}^gT \leq S$. Therefore ${}^gT \leq Z_S({}^gP; X) = Z({}^gP; X)$ by the \mathcal{F} -centricity at X of P and we conclude that $T = Z(P; X)$, as desired. \square

The main purpose of introducing the abstract fusion-centric way of thinking of X -centricity at this point is that it makes certain results cleaner to prove:

Proposition 3.1.6. If X and Y are \mathcal{F} -stable S -sets and there is a surjective map of S -sets $f : X \rightarrow Y$, then every Y -centric subgroup is also X -centric.

Proof. Let C and D be the cores of X and Y , respectively. The existence of such a surjection forces $C \leq D$. Thus $Z_S(P) \cap C \leq Z_S(P) \cap D$, so if $Z_S(P) \cap D = Z(P) \cap D$ then $Z_S(P) \cap C = Z(P) \cap C$. The result follows. \square

Corollary 3.1.7. p -centricity implies X -centricity for all \mathcal{F} -stable X .

In particular, every p - or \mathcal{F} -centric (both henceforth “centric”) subgroup of S is automatically X -centric for every \mathcal{F} -stable X . Moreover, the more faithful X is (so the smaller the core C is), the easier it is for subgroups of S to be X -centric, to the point where if X is a faithful S -set, every $P \leq S$ is X -centric.

Definition 3.1.8. Let \mathcal{T}_G^{cX} denote the full subcategory of \mathcal{T}_G whose objects are the X -centric subgroups. Similarly, for \mathfrak{X} an abstract action with finite set X , let \mathfrak{X}^{cX} be the full subcategory on the X -centric subgroups.

The first thing we must show is that we have not lost too much information by this restriction. Let $\mathbb{X} : \mathcal{T}_G^{cX} \rightarrow \mathcal{TOP}$ be the functor

$$\begin{array}{ccc} P & \dashrightarrow & X \\ g \downarrow & & \downarrow \ell_g \\ Q & \dashrightarrow & X \end{array}$$

Proposition 3.1.9. There exists a mod- p equivalence $\text{hocolim}_{\mathcal{T}_G^{cX}} \mathbb{X} \simeq_p B_G X$.

Proof. Let $\mathcal{E}G$ be the contractible groupoid (meaning there is precisely one morphism from any object to any other object) with objects $g \in G$, \mathcal{K}_{cX} the category associated to the poset of X -centric subgroups, and K_{cX} the associated G -simplicial complex. We view the G -set X as a discrete G -category. If $\mathcal{G}(\mathbb{X})$ is the Grothendieck construction associated to the functor \mathbb{X} , we have $|\mathcal{G}(\mathbb{X})| \simeq \operatorname{hocolim}_{\mathcal{T}_G^{cX}} \mathbb{X}$ by [Tho], so we will use the Grothendieck construction as our model for the homotopy colimit.

There is a functor $\mathcal{E}G \times_G (X \times \mathcal{K}_{cX}) \rightarrow \mathcal{G}(\mathbb{X})$ given by

$$\begin{array}{ccc} [g, (x, P)] & \dashrightarrow & ({}^gP, g \cdot x) \\ \downarrow [g \mapsto hg, (id_x, P \leq Q)] & & \downarrow h \\ [hg, (x, Q)] & \dashrightarrow & ({}^{hg}Q, hg \cdot x) \end{array}$$

that is easily seen to have an inverse isomorphism of categories

$$\begin{array}{ccc} (P, x) & \dashrightarrow & [1, (x, P)] \\ \downarrow h & & \downarrow [1 \mapsto h, (id_x, P \leq h^{-1}Q)] \\ (Q, h \cdot x) & \dashrightarrow & [1, (h \cdot x, Q)] = [h, (x, h^{-1}Q)] \end{array}$$

Thus, on taking realizations, we get a homeomorphism of spaces

$$EG \times_G (X \times K_{cX}) \simeq \operatorname{hocolim}_{\mathcal{T}_G^{cX}} \mathbb{X}.$$

Finally, consider the natural projection map $EG \times_G (X \times K_{cX}) \rightarrow EG \times_G X$; if we can show that this is a mod- p homology isomorphism, the result will follow.

By Corollary 3.1.7, the collection of X -centric subgroups contains all centric subgroups, and it is easy to check that it is closed under p -overgroups. Thus [Dwy, Theorem 8.3] applies to show that K_{cX} is \mathbb{F}_p -acyclic, and therefore the natural map $EG \times_G (X \times K_{cX}) \rightarrow EG \times_G X$, being a fibration with fiber K_{cX} , is a mod- p homology isomorphism by the Serre spectral sequence. \square

There is a more drastic way of reducing information in \mathcal{T}_G than simply restricting our attention to various subgroups, in which we quotient out p' -information directly.

Note that there is a free right action of $Z(P; X)$ on $\mathcal{T}_G(P, Q)$, so we can consider categories whose morphisms are orbits of subgroups of $Z(P; X)$ in the transporter system. The need to lose only p' -information is one reason why we have restricted our attention to the X -centric subgroups of S , as we shall now see.

Recall that for H a finite group, $O^p(H)$ is the smallest normal subgroup of H of p -power index. If P is X -centric, $Z_G(P; X) = Z(P; X) \times Z'_G(P; X)$ for $Z'_G(P; X)$ a p' -group, and we have $O^p(Z_G(P; X)) = Z'_G(P; X)$.

Definition 3.1.10. For a G -set X , define the p -centric linking action system of S at X to be the category $\mathcal{L}_G^{cX} = \mathcal{L}_S^{cX}(G)$ whose objects are X -centric subgroups of S , and where $\mathcal{L}_G^{cX}(P, Q) = N_G(P, Q)/O^p(Z_G(P; X))$.

Just as we formed the Borel construction $B_G X$ up to p -completion by considering the homotopy colimit of a functor $\mathcal{T}_G^{cX} \rightarrow \mathcal{TOP}$, there is a similarly defined functor whose homotopy colimit is of particular interest to us in the context of linking action systems. Let $\overline{\mathbb{X}}: \mathcal{L}_G^{cX} \rightarrow \mathcal{TOP}$ be the functor

$$\begin{array}{ccc} P & \dashrightarrow & X \\ \downarrow [g] & & \downarrow \ell_g \\ Q & \dashrightarrow & X \end{array}$$

By construction, the natural quotient $\pi: \mathcal{T}_G^{cX} \rightarrow \mathcal{L}_G^{cX}$ satisfies the conditions of [BLO1, Lemma 1.3]. In particular, this implies

Proposition 3.1.11. *The quotient $\pi : \mathcal{T}_G^{cX} \rightarrow \mathcal{L}_G^{cX}$ induces a mod- p equivalence:*

$$\mathrm{hocolim}_{\mathcal{L}_G^{cX}} \overline{\mathbb{X}} \simeq_p \mathrm{hocolim}_{\mathcal{T}_G^{cX}} \overline{\mathbb{X}} \circ \pi.$$

Theorem 3.1.12.

$$\mathrm{hocolim}_{\mathcal{L}_G^{cX}} \mathbb{X} \simeq_p B_G X$$

Proof. Combine Propositions 3.1.9 and 3.1.11 with the fact that $\mathbb{X} = \overline{\mathbb{X}} \circ \pi$. □

3.2 Classifying spaces of abstract fusion actions

3.2.1 Stating the problem

In this section let \mathfrak{X} be a fixed abstract fusion action system with underlying fusion system \mathcal{F} . The goal is to describe what a “classifying space” for \mathfrak{X} should look like. The heuristic is that $B\mathfrak{X}$ should recover the homotopy type of the Borel construction of X , up to p -completion, just as in the ambient case.

The problem with the heuristic is that, without an ambient group, we lack a Borel construction to aim for. Section 3.1 tells us that we could perhaps instead try to define $B\mathfrak{X}$ to be the homotopy colimit of some functor into \mathcal{TOP} , but again without an ambient group we do not know what the source category for such a functor should look like. We could try to rectify the situation by developing some notion of an *abstract linking action system*. This would be a category $\mathcal{L}^{\mathfrak{X}}$ associated to \mathfrak{X} that would have the “right” properties so that we could construct a functor $\mathbb{X} : \mathcal{L}^{\mathfrak{X}} \rightarrow \mathcal{TOP}$, all of which would be a generalization of the ambient case of Section 3.1.

This will in fact be the plan of attack take, but first a detour: In order to properly understand what is meant by $\mathcal{L}^{\mathfrak{X}}$ having the “right” properties, we should try to understand the space we are looking for purely in terms that can be described by the fusion action system itself. The question then becomes what spaces we can make from the category \mathfrak{X} . The first guess of $|\mathfrak{X}|$ will not give us what we want, as this relies only on the shape of the category of \mathfrak{X} and does not take into account the fact that it should be thought of as a combination of a diagram in groups together with permutations of X .

More accurately, we should think of \mathfrak{X} as a diagram in groupoids via the functor $\mathcal{B}_-X : \mathfrak{X} \rightarrow \mathcal{GRP}\mathcal{D}$. For $P \leq S$, let $\mathcal{B}_P X$ denote the translation category of the P -set X . Recall that the objects of $\mathcal{B}_P X$ are the points of X and the morphisms are given by $\mathcal{B}_P X(x, x') = \{\vec{p} | p \in P \text{ and } p \cdot x = x'\}$. We then define \mathcal{B}_-X to be the functor

$$\begin{array}{ccc} P & \dashrightarrow & \mathcal{B}_P X \\ (\varphi, \sigma) \downarrow & & \downarrow \mathcal{B}(\varphi, \sigma) \\ Q & \dashrightarrow & \mathcal{B}_Q X \end{array}$$

where $\mathcal{B}(\varphi, \sigma)$ is the functor that acts on objects by σ and morphisms by φ .

If we then set $B_-X : \mathfrak{X} \rightarrow \mathcal{TOP}$ to be $|\mathcal{B}_-X|$, we can consider the space $\mathrm{hocolim}_{\mathfrak{X}} B_-X$, which is defined with information contained in \mathfrak{X} .

Unfortunately, this space is not what we want if S is nonabelian or acts nontrivially on X , as the following example illustrates:

Consider the case that $X = *$ is the trivial S -set. In this case $\mathfrak{X} = \mathcal{F}$ and the space in question is $\mathrm{hocolim}_{\mathcal{F}} BP$, while the space we want to recover up to p -completion is $EG \times_G * = BG$. In particular, the natural map $BS \rightarrow BG$ is Sylow, which is to say if R is any p -group and we have a map $BR \rightarrow BG$, there is a factorization

$$\begin{array}{ccc} BS & \xrightarrow{\quad} & BG \\ & \nwarrow \text{dashed} \nearrow & \\ & BR & \end{array}$$

up to homotopy. If S is nonabelian, $\text{Inn}(S)$ is nontrivial, and therefore $R = \text{Inn}(S) \ltimes S$ contains S as a proper subgroup. Let $\mathcal{G} = \mathcal{G}(\mathcal{B}-)$ be the Grothendieck category for the functor $\mathcal{B} : \mathcal{F} \rightarrow \mathcal{GRP}$, so that

$$\text{hocolim}_{\mathcal{F}} \mathcal{B}- = |\mathcal{G}|$$

By definition $\mathcal{G}(S) = \text{Inn}(S) \ltimes S$, so there is an obvious map of spaces $BR \rightarrow |\mathcal{G}|$ that does not factor through BS .

The problem with simply taking $\text{hocolim}_{\mathfrak{X}} \mathcal{B}-X$ is that this construction is in some sense double-counting the noncentral elements of S , as well as the elements not in K . In particular, any non- X -central $s \in S$ defines both a morphism in $\mathfrak{s} \in \mathcal{B}S$ and a morphism $(c_s, \ell_s) \in \mathfrak{X}(S)$; by simply taking the homotopy colimit of $\mathcal{B}-X$ these separate morphisms both contribute even though they come from the same element of S . In other words, $\text{hocolim}_{\mathfrak{X}} \mathcal{B}-X$ is too big, in that it has too many arrows.

So let's kill the offending morphisms.

Definition 3.2.1. The *orbit category* of \mathfrak{X} is the category $\mathcal{O}^{\mathfrak{X}} := \mathcal{O}(\mathfrak{X})$ whose objects are the subgroups of S and whose morphisms are given by $\mathcal{O}^{\mathfrak{X}}(P, Q) = Q \backslash \mathfrak{X}(P, Q)$. In other words, the Hom-set from P to Q is the set orbits of the Q -action of $\mathfrak{X}(P, Q)$ given by postcomposition by (c_q, ℓ_q) .

The full subcategory of $\mathcal{O}^{\mathfrak{X}}$ whose objects are the \mathcal{F} -centric subgroups of S will be denoted $\mathcal{O}^{c\mathfrak{X}}$.

The functor $\mathcal{B}-X$ does not descend to $\mathcal{O}^{c\mathfrak{X}} \rightarrow \mathcal{TOP}$, but because $\mathcal{O}^{c\mathfrak{X}}$ is defined by quotienting out inner automorphisms, it is easy to see that there is a homotopy functor $\overline{\mathcal{B}}-X : \mathcal{O}^{c\mathfrak{X}} \rightarrow \text{ho}\mathcal{TOP}$. If we could find a homotopy lifting $\tilde{\mathcal{B}}$ as in Figure 4 we could consider the space $\text{hocolim}_{\mathcal{O}^{c\mathfrak{X}}} \tilde{\mathcal{B}}-X$ as a possibility for the classifying space of \mathcal{F}^X .

$$\begin{array}{ccc} & & \mathcal{TOP} \\ & \nearrow \tilde{\mathcal{B}}-X & \downarrow \\ \mathcal{O}^{c\mathfrak{X}} & \xrightarrow{\overline{\mathcal{B}}-X} & \text{ho}\mathcal{TOP} \end{array}$$

Figure 4: The homotopy lifting problem

We shall see that the homotopy colimit of such a lifting $\tilde{\mathcal{B}}$ is, in fact, the solution we have been seeking.

Proposition 3.2.2. *In the presence of an ambient group G acting on X and giving rise to \mathfrak{X} ,*

$$\text{hocolim}_{\mathcal{O}^{c\mathfrak{X}}} \tilde{\mathcal{B}}-X \simeq_p B_G X.$$

Proof. Follows immediately from Problem 3.2.3, whose solution is detailed in Subsection 3.2.2. \square

We are now in the position to fully state the problem we want to solve:

Problem 3.2.3. Given an abstract fusion action system \mathfrak{X} , construct an *abstract linking action system* associated to \mathfrak{X} . This should be a category $\mathcal{L}^{\mathfrak{X}}$ together with functors

$$\pi : \mathcal{L}^{\mathfrak{X}} \rightarrow \mathfrak{X}^{cX} \quad \text{and} \quad \mathbb{X} : \mathcal{L}^{\mathfrak{X}} \rightarrow \mathcal{TOP}$$

such that

- In the presence of an ambient group G , this \mathcal{L}_G^{cX} and the functor $\overline{\mathbb{X}}$ described in Subsection 3.1 give an abstract linking action system associated to \mathfrak{X}_G .
- If $\overline{\pi} : \mathcal{L}^{\mathfrak{X}} \rightarrow \mathcal{O}^{c\mathfrak{X}}$ is the composite $\mathcal{L}^{\mathfrak{X}} \xrightarrow{\pi} \mathfrak{X}^{cX} \rightarrow \mathcal{O}^{c\mathfrak{X}}$, then the left homotopy Kan extension of \mathbb{X} over $\overline{\pi}$ is a homotopy lifting of $\overline{\mathcal{B}}-X : \mathcal{O}^{c\mathfrak{X}} \rightarrow \text{ho}\mathcal{TOP}$, which will be denoted $\tilde{\mathcal{B}}_{\mathcal{L}^{\mathfrak{X}}} : \mathcal{O}^{c\mathfrak{X}} \rightarrow \mathcal{TOP}$. Consequently,

$$\text{hocolim}_{\mathcal{L}^{\mathfrak{X}}} \mathbb{X} \simeq \text{hocolim}_{\mathcal{O}^{c\mathfrak{X}}} \tilde{\mathcal{B}}_{\mathcal{L}^{\mathfrak{X}}}.$$

3.2.2 Solving the problem

Definition 3.2.4. An *(X-centric) linking action system* associated to the fusion action system \mathfrak{X} is a category $\mathcal{L}^{\mathfrak{X}}$ whose objects are the X -centric subgroups of S , together with a pair of functors

$$\mathcal{T}_S^{cX} \xrightarrow{\delta} \mathcal{L}^{\mathfrak{X}} \xrightarrow{\pi} \mathfrak{X}^{cX}.$$

For each $s \in N_S(P, Q)$, let \widehat{s} be the corresponding morphism $\delta_{P, Q}(s) \in \mathcal{L}^{\mathfrak{X}}(P, Q)$. For $\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(P, Q)$ denote the components of $\pi_{P, Q}(\mathfrak{g})$ by the pair $(c_{\mathfrak{g}}, \ell_{\mathfrak{g}})$.

The following conditions must be satisfied:

- (A) The functors δ and π are the identity on objects. δ is injective and π is surjective on morphisms. Moreover, $Z(P; X)$ acts (via δ) right-freely on $\mathcal{L}^{\mathfrak{X}}(P, Q)$, and $\pi_{P, Q} : \mathcal{L}^{\mathfrak{X}}(P, Q) \rightarrow \mathfrak{X}^{cX}(P, Q)$ is the orbit map of this action.
- (B) For each $p \in P$, we have $c_{\widehat{p}} = c_p : P \rightarrow P$ and $\ell_{\widehat{p}} = \ell_p : X \rightarrow X$.
- (C) For each $\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(P, Q)$ and $p \in P$, the following commutes in $\mathcal{L}^{\mathfrak{X}}$:

$$\begin{array}{ccc} P & \xrightarrow{\mathfrak{g}} & Q \\ \widehat{p} \downarrow & & \downarrow \widehat{c_{\mathfrak{g}}(p)} \\ P & \xrightarrow{\mathfrak{g}} & Q \end{array}$$

Finally, every linking action system comes naturally equipped with a functor $\mathbb{X} : \mathcal{L}^{\mathfrak{X}} \rightarrow \mathcal{TOP}$ defined by

$$\begin{array}{ccc} P & \dashrightarrow & X \\ \mathfrak{g} \downarrow & & \downarrow \ell_{\mathfrak{g}} \\ Q & \dashrightarrow & X \end{array}$$

where X is regarded as a discrete topological space.

Remark 3.2.5. It is easy to see that:

- If $X = *$ is the trivial S -set and $\mathfrak{X} = \mathcal{F}$ is a saturated fusion system, then an abstract linking action system is the same as a centric linking system as defined in [BLO2]. In this case \mathbb{X} is just the trivial functor $\mathcal{L}^* \rightarrow \mathcal{TOP}$ whose homotopy colimit is $|\mathcal{L}^*|$.
- In the presence of an ambient group G , the category \mathcal{L}_G^{cX} of Subsection 3.1 is an example of an abstract linking action system associated to \mathfrak{X}_G . In particular, this definition satisfies the first condition of Problem 3.2.3.

Unless it is necessary to emphasize the fact that we are looking at the X -centric subcategories, we shall omit explicit notational reference to them.

The remainder of this section is devoted to proving that these abstract linking action systems satisfy the second condition of Problem 3.2.3. We begin with some basic properties of abstract linking action systems:

Proposition 3.2.6. *Let $P \xrightarrow{(\varphi, \sigma)} Q \xrightarrow{(\psi, \tau)} R$ be a sequence of morphisms in \mathfrak{X} . Then for any*

$$\mathfrak{g} \in \pi_{Q, R}^{-1}((\psi, \tau)) \subseteq \mathcal{L}^{\mathfrak{X}}(Q, R) \quad \text{and} \quad \widetilde{\mathfrak{gh}} \in \pi_{P, R}^{-1}((\psi\varphi, \tau\sigma)) \subseteq \mathcal{L}^{\mathfrak{X}}(P, R)$$

there is a unique $\mathfrak{h} \in \pi_{P, Q}^{-1}((\varphi, \sigma)) \subseteq \mathcal{L}^{\mathfrak{X}}(P, Q)$ such that $\mathfrak{gh} = \widetilde{\mathfrak{gh}}$.

Proof. Pick any $\mathfrak{h}' \in \pi_{P,Q}^{-1}((\varphi, \sigma)) \subseteq \mathcal{L}^{\mathfrak{X}}(P, Q)$. By Axiom (A), the X -center $Z(P; X)$ acts freely and transitively on $\pi_{P,R}^{-1}((\psi\varphi, \tau\sigma)) \subseteq \mathcal{L}^{\mathfrak{X}}(P, R)$, so there is a unique $z \in Z(P; X)$ such that $\widetilde{\mathfrak{g}\mathfrak{h}} = \mathfrak{g}\mathfrak{h}'\widehat{z}$. Therefore setting $\mathfrak{h} = \mathfrak{h}'\widehat{z}$ gives a morphism in $\pi_{P,Q}^{-1}((\varphi, \sigma))$ such that $\mathfrak{g}\mathfrak{h} = \widetilde{\mathfrak{g}\mathfrak{h}}$.

To prove uniqueness, suppose that we have $\mathfrak{h}_1, \mathfrak{h}_2 \in \mathcal{L}^{\mathfrak{X}}(P, Q)$ such that $\mathfrak{g}\mathfrak{h}_1 = \mathfrak{g}\mathfrak{h}_2$, \mathfrak{h}_i lifts (φ, ψ) , and $\mathfrak{g}\mathfrak{h}_i = \widetilde{\mathfrak{g}\mathfrak{h}}$. Then $c_{\mathfrak{g}}c_{\mathfrak{h}_1} = c_{\mathfrak{g}}c_{\mathfrak{h}_2} \in \mathfrak{X}(P, R)$, and since every morphism of a fusion action system is mono (the first coordinate is an injective group map and the second coordinate is invertible), we conclude that $c_{\mathfrak{h}_1} = c_{\mathfrak{h}_2}$. Therefore by Axiom (A) there is a unique $z \in Z(P; X)$ so that $\mathfrak{h}_1 = \mathfrak{h}_2 \circ \widehat{z}$. Then the facts that

$$\widetilde{\mathfrak{g}\mathfrak{h}} = \mathfrak{g}\mathfrak{h}_1 = \mathfrak{g}\mathfrak{h}_2 \circ \widehat{z} = \widetilde{\mathfrak{g}\mathfrak{h}} \circ \widehat{z}$$

and $Z(P; X)$ acts right-freely on $\mathcal{L}^{\mathfrak{X}}(P, R)$ implies that $z = 1$, so $\mathfrak{h}_1 = \mathfrak{h}_2$. \square

Remark 3.2.7. Note that the full strength of the Axioms of a linking action system is not used to prove this proposition. In particular, the fact that $\mathcal{L}^{\mathfrak{X}}$ comes equipped with a functor $\mathcal{T}_S^{cX} \rightarrow \mathcal{L}^{\mathfrak{X}}$ could have been replaced with the (apparently) weaker assumption that for every X -centric $P \leq S$, there is an injection $P \rightarrow \mathcal{L}^{\mathfrak{X}}(P)$ such that Axiom (A) holds. This observation will sometimes be useful when trying to construct linking action systems.

Proposition 3.2.6 has a number of immediate consequences:

Corollary 3.2.8. *Every morphism in $\mathcal{L}^{\mathfrak{X}}$ is categorically mono.*

Proof. Suppose that we're given morphisms $\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(Q, R)$ and $\mathfrak{h}, \mathfrak{h}' \in \mathcal{L}^{\mathfrak{X}}(P, Q)$ such that $\mathfrak{g}\mathfrak{h} = \mathfrak{g}\mathfrak{h}'$. Then $\pi(\mathfrak{g}\mathfrak{h}) = \pi(\mathfrak{g}\mathfrak{h}')$, and \mathfrak{g} is a lifting of $\pi(\mathfrak{g})$, so the uniqueness statement of Proposition 3.2.6 forces $\mathfrak{h} = \mathfrak{h}'$. \square

Corollary 3.2.9. *If $\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(P, Q)$ is such that $(c_{\mathfrak{g}}, \ell_{\mathfrak{g}}) \in \mathfrak{X}(P, Q)$ is an isomorphism, then \mathfrak{g} is itself an isomorphism.*

Proof. Apply Proposition 3.2.6 to the sequence $Q \xrightarrow{(c_{\mathfrak{g}}, \ell_{\mathfrak{g}})^{-1}} P \xrightarrow{(c_{\mathfrak{g}}, \ell_{\mathfrak{g}})} Q$ with \mathfrak{g} lifting $(c_{\mathfrak{g}}, \ell_{\mathfrak{g}})$ and $\text{id}_Q^{\mathcal{L}^{\mathfrak{X}}}$ lifting the composite $\text{id}_Q^{\mathfrak{X}}$. We obtain a unique $\mathfrak{h} \in \mathcal{L}^{\mathfrak{X}}(Q, P)$ such that $\mathfrak{g}\mathfrak{h} = \text{id}_Q^{\mathcal{L}^{\mathfrak{X}}}$. It follows that $\mathfrak{g}\mathfrak{h}\mathfrak{g} = \mathfrak{g} = \mathfrak{g} \circ \text{id}_P^{\mathcal{L}^{\mathfrak{X}}}$ and thus $\mathfrak{h}\mathfrak{g} = \text{id}_P^{\mathcal{L}^{\mathfrak{X}}}$ as \mathfrak{g} is categorically mono. Therefore $\mathfrak{h} = \mathfrak{g}^{-1}$. \square

Notation 3.2.10. We denote by i_P^Q the morphism $\delta_{P,Q}(1) \in \mathcal{L}^{\mathfrak{X}}(P, Q)$, and call this the “inclusion” of P in Q in $\mathcal{L}^{\mathfrak{X}}$.

Corollary 3.2.11. *For any $\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(P, Q)$ and X -centric subgroups $P^* \leq P$ and $Q^* \leq Q$ such that $c_{\mathfrak{g}}(P^*) \leq Q^*$, there is a unique morphism $\text{res}_{P^*}^{Q^*}(\mathfrak{g}) \in \mathcal{L}^{\mathfrak{X}}(P^*, Q^*)$ such that the following diagram commutes in $\mathcal{L}^{\mathfrak{X}}$:*

$$\begin{array}{ccc} P & \xrightarrow{\mathfrak{g}} & Q \\ i_{P^*}^P \uparrow & & \uparrow i_{Q^*}^Q \\ P^* & \xrightarrow{\text{res}_{P^*}^{Q^*}(\mathfrak{g})} & Q^* \end{array}$$

Proof. Apply Proposition 3.2.6 to the diagram $P^* \xrightarrow{c_{\mathfrak{g}}|_{P^*}} Q^* \xrightarrow{i_{Q^*}^Q} Q$ with $i_{Q^*}^Q$ lifting $i_{Q^*}^Q$ and $\mathfrak{g} \circ i_{P^*}^P$ lifting the composite. \square

Notation 3.2.12. The morphism $\text{res}_{P^*}^{Q^*}(\mathfrak{g})$ is called the “restriction” of \mathfrak{g} , and will sometimes be denoted $\mathfrak{g}|_{P^*}^{Q^*}$ or just $\mathfrak{g}|_{P^*}$ if the target is clear.

Corollary 3.2.13. *Every morphism in $\mathcal{L}^{\mathfrak{X}}$ factors uniquely as an isomorphism followed by an inclusion.*

Proof. For $\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(P, Q)$, the morphism $(c_{\mathfrak{g}}, \ell_{\mathfrak{g}}) \in \mathfrak{X}(P, c_{\mathfrak{g}}(P))$ is an isomorphism in the underlying action system, so $\text{res}_P^{c_{\mathfrak{g}}(P)}(\mathfrak{g})$ is an isomorphism in $\mathcal{L}^{\mathfrak{X}}$ by Corollary 3.2.9. Clearly $\mathfrak{g} = \mathfrak{i}_{c_{\mathfrak{g}}(P)}^Q \circ \text{res}_P^{c_{\mathfrak{g}}(P)}(\mathfrak{g})$, and uniqueness now follows by another application of Proposition 3.2.6. \square

Proposition 3.2.14. *Every morphism in $\mathcal{L}^{\mathfrak{X}}$ is categorically epi.*

Proof. By Corollary 3.2.13, it suffices to show that $\mathfrak{i} := \mathfrak{i}_P^Q$ is epi for all X -centric $P \leq Q$. Moreover, it suffices to assume that $P \trianglelefteq Q$, as any inclusion of p -groups can be refined to a sequence of normal inclusions. Let $\mathfrak{g}_1, \mathfrak{g}_2 \in \mathcal{L}^{\mathfrak{X}}(Q, R)$ be two morphisms such that $\mathfrak{g}_1 \circ \mathfrak{i} = \mathfrak{g}_2 \circ \mathfrak{i}$; we want to show that $\mathfrak{g}_1 = \mathfrak{g}_2$.

The image of \mathfrak{g}_i in $\mathfrak{X}(Q, R)$ is $(c_{\mathfrak{g}_i}, \sigma_i)$ and the assumption that $\mathfrak{g}_1 \circ \mathfrak{i} = \mathfrak{g}_2 \circ \mathfrak{i}$ implies that $\sigma_1 = \sigma_2$. We first show that it also follows that $c_{\mathfrak{g}_1} = c_{\mathfrak{g}_2}$.

Note that the assumption that $P \trianglelefteq Q$ implies that every $q \in Q$ determines a morphism in $\widehat{q} \in \mathcal{L}^{\mathfrak{X}}(P)$, which is easily seen to be the restriction of $\widehat{q} \in \mathcal{L}^{\mathfrak{X}}(Q)$. Then the diagram

$$\begin{array}{ccc} Q & \xrightarrow{\mathfrak{g}_i} & R \\ \widehat{q} \downarrow & & \downarrow \widehat{c_{\mathfrak{g}_i}(q)} \\ Q & \xrightarrow{\mathfrak{g}_i} & R \end{array}$$

commutes in $\mathcal{L}^{\mathfrak{X}}$ by Axiom (C) for $i = 1, 2$. The assumption on the \mathfrak{g}_i also implies that their restrictions to P in $\mathcal{L}^{\mathfrak{X}}$ are equal; denote this common morphism by $\mathfrak{h} \in \mathcal{L}^{\mathfrak{X}}(P, c_{\mathfrak{h}}(P))$. We can therefore form the restriction of this diagram, which gives us

$$\begin{array}{ccc} P & \xrightarrow{\mathfrak{h}} & c_{\mathfrak{h}}(P) \\ \widehat{q} \downarrow & & \downarrow \widehat{c_{\mathfrak{g}_i}(q)} \\ P & \xrightarrow{\mathfrak{h}} & c_{\mathfrak{g}}(P) \end{array}$$

for $i = 1, 2$. As all the morphisms in this restriction diagram are iso, and three of the four do not depend on choice of i , we conclude that $\widehat{c_{\mathfrak{g}_1}(q)} = \widehat{c_{\mathfrak{g}_2}(q)}$ for all $q \in Q$. The assignment $q \mapsto \widehat{q}$ is injective, so we conclude that $c_{\mathfrak{g}_1} = c_{\mathfrak{g}_2}$ on Q .

Thus we have $\pi(\mathfrak{g}_1) = \pi(\mathfrak{g}_2)$, so by Axiom (2), there is some $z \in Z(Q; X)$ such that $\mathfrak{g}_2 = \mathfrak{g}_1 \circ \widehat{z}$. We then compute $\mathfrak{g}_1 \circ \widehat{z} \circ \mathfrak{i} = \mathfrak{g}_2 \circ \mathfrak{i} = \mathfrak{g}_1 \circ \mathfrak{i}$. The fact that \mathfrak{g}_1 is mono implies that $\widehat{z} \circ \mathfrak{i} = \mathfrak{i}$.

Finally, the fact that $P \leq Q$ implies that $z \in Z(P; X)$, and Axiom (3) again shows that $\widehat{z} \circ \mathfrak{i} = \mathfrak{i} \circ \widehat{z}$ (for \widehat{z} respectively an isomorphism in $\mathcal{L}^{\mathfrak{X}}$ of Q and P). The freeness of the right action of $Z(P; X)$ on $\mathcal{L}^{\mathfrak{X}}(P, Q)$ forces $z = 1$, and the result is proved. \square

Corollary 3.2.15. *Extensions of morphisms in $\mathcal{L}^{\mathfrak{X}}$ are unique when they exist. In other words, for any $\mathfrak{g}^* \in \mathcal{L}^{\mathfrak{X}}(P^*, Q^*)$ and subgroups $P^* \leq P$ and $Q^* \leq Q$, there is at most one morphism $\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(P, Q)$ such that the diagram*

$$\begin{array}{ccc} P & \xrightarrow{\mathfrak{g}} & Q \\ \mathfrak{i}_{P^*}^P \uparrow & & \uparrow \mathfrak{i}_{Q^*}^Q \\ P^* & \xrightarrow{\mathfrak{g}^*} & Q^* \end{array}$$

commutes in $\mathcal{L}^{\mathfrak{X}}$.

Proof. If \mathfrak{g}' is another such extension, the equalities $\mathfrak{g} \circ \mathfrak{i}_{P^*}^P = \mathfrak{i}_{Q^*}^Q \circ \mathfrak{g}^* = \mathfrak{g}' \circ \mathfrak{i}_{P^*}^P$ and the fact that $\mathfrak{i}_{P^*}^P$ is epi force $\mathfrak{g} = \mathfrak{g}'$. \square

Proposition 3.2.16. *If $P \leq S$ is fully normalized in \mathcal{F} , then $\widehat{N_S(P)}|_P^P \in \text{Syl}_p(\mathcal{L}^{\mathfrak{X}}(P))$.*

Proof. By Axiom (A) of linking action systems, $|\mathcal{L}^{\mathfrak{X}}(P)| = |\mathfrak{X}(P)| \cdot |Z(P; X)|$ and by definition of $\text{Aut}_S(P; X)$,

$$|N_S(P)| = |\text{Aut}_S(P; X)| \cdot |Z_S(P; X)| = |\text{Aut}_S(P; X)| \cdot |Z(P; X)|.$$

As P is fully normalized, the Saturation Axioms imply that $\text{Aut}_S(P; X) \in \text{Syl}_p(\mathfrak{X}(P))$, and the result easily follows. \square

Notation 3.2.17. Let $\bar{\pi}$ be the composite $\mathcal{L}^{\mathfrak{X}} \xrightarrow{\pi} \mathfrak{X} \rightarrow \mathcal{O}^{c\mathfrak{X}}$.

Proposition 3.2.18. *Suppose that $\mathfrak{g}, \mathfrak{h} \in \mathcal{L}^{\mathfrak{X}}(P, Q)$ are such that $\bar{\pi}(\mathfrak{g}) = \bar{\pi}(\mathfrak{h})$. Then there is a unique element $q \in Q$ such that $\mathfrak{h} = \hat{q} \circ \mathfrak{g}$. In other words, the map $\bar{\pi}_{P, Q} : \mathcal{L}^{\mathfrak{X}}(P, Q) \rightarrow \mathcal{O}^{c\mathfrak{X}}(P, Q)$ is the orbit map of the free left action of Q on $\mathcal{L}^{\mathfrak{X}}(P, Q)$.*

Proof. The condition on \mathfrak{g} and \mathfrak{h} implies that there is some $q' \in Q$ such that

$$(c_{\mathfrak{h}}, \ell_{\mathfrak{h}}) = (c_{q'}, \ell_{q'}) \circ (c_{\mathfrak{g}}, \ell_{\mathfrak{g}}).$$

Condition (A) then implies that there is a unique $z \in Z(P; X)$ such that

$$\mathfrak{h} = \hat{q}' \circ \mathfrak{g} \circ \hat{z} = \hat{q}' \circ \widehat{c_{\mathfrak{g}}(z)} \circ \mathfrak{g}$$

where the second equality follows from Condition (C). Setting $q = q' \circ c_{\mathfrak{g}}(z)$ gives the desired element.

The uniqueness of q is a direct consequence of the fact that \mathfrak{g} is epi and the assignment $q \mapsto \hat{q}$ is injective. \square

We now have the necessary results to show that we have solved the second point of Problem 3.2.3. Recall that $\tilde{B}_{-X} : \mathcal{O}^{c\mathfrak{X}} \rightarrow \mathcal{TOP}$ is the left homotopy Kan extension of $\mathbb{X} : \mathcal{L}^{\mathfrak{X}} \rightarrow \mathcal{TOP}$ over $\bar{\pi} : \mathcal{L}^{\mathfrak{X}} \rightarrow \mathcal{O}^{c\mathfrak{X}}$, and that we are looking for a homotopy lifting of $\tilde{B}_{-X} : \mathcal{O}^{c\mathfrak{X}} \rightarrow \mathcal{TOP}$, as depicted in Figure 5.

$$\begin{array}{ccc} \mathcal{L}^{\mathfrak{X}} & \xrightarrow{\mathbb{X}} & \mathcal{TOP} \\ \bar{\pi} \downarrow & \nearrow \tilde{B}_{-X} & \\ \mathcal{O}^{c\mathfrak{X}} & & \end{array} \qquad \begin{array}{ccc} & & \mathcal{TOP} \\ & \nearrow ? & \downarrow \\ \mathcal{O}^{c\mathfrak{X}} & \xrightarrow{\tilde{B}_{-X}} & \mathcal{TOP} \end{array}$$

Figure 5: What we have vs. What we want

For ease of notation, write \tilde{B} for \tilde{B}_{-X} and B for \overline{B}_{-X} .

Proposition 3.2.19. *The functor \tilde{B} is a homotopy lifting of B .*

Proof. On objects: We want to show that $\tilde{B}(P) \simeq B(P) = B_P X$. Recall that

$$\tilde{B}(P) \simeq \text{hocolim}_{(Q, \alpha) \in (\bar{\pi} \downarrow P)} \mathbb{X}(Q)$$

where $(\bar{\pi} \downarrow P)$ is the overcategory of $\bar{\pi}$ over P . The objects are pairs (Q, α) , where Q is an X -centric subgroup of S , viewed as either an object of $\mathcal{L}^{\mathfrak{X}}$ or $\mathcal{O}^{c\mathfrak{X}}$ as appropriate, and α is a morphism in $\mathcal{O}^{c\mathfrak{X}}(Q, P)$. A morphism from (Q, α) to (R, β) is $\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(Q, R)$ such that the following commutes in $\mathcal{O}^{c\mathfrak{X}}$:

$$\begin{array}{ccc} Q & \xrightarrow{\bar{\pi}(\mathfrak{g})} & R \\ & \searrow \alpha & \swarrow \beta \\ & P & \end{array}$$

Thomason's theorem [Tho] shows that this homotopy colimit expression of $\overline{\mathcal{B}}(P)$ has the homotopy type of $|\mathcal{G}(P)|$, where $\mathcal{G}(P)$ is the Grothendieck category associated to $(\overline{\pi} \downarrow P)$ and the functor \mathbb{X} .

Explicitly, $\mathcal{G}(P)$ is the category whose objects are pairs $((Q, \alpha), x)$ where (Q, α) is an object of $(\overline{\pi} \downarrow P)$ and $x \in X$. A morphism $((Q, \alpha), x) \rightarrow ((R, \beta), y)$ is a morphism $\mathbf{g} \in (\overline{\pi} \downarrow P)((Q, \alpha), (R, \beta))$ —and therefore $\mathbf{g} \in \mathcal{L}^{\mathfrak{X}}(Q, R)$ —such that $\ell_{\mathbf{g}}(x) = y$.

Let $\check{\mathcal{B}}(P) \subseteq \mathcal{G}(P)$ be the subcategory whose objects are pairs $((P, \text{id}_P), x)$ for all $x \in X$, and where

$$\check{\mathcal{B}}(P)((P, \text{id}_P), x), ((P, \text{id}_P), y)) = \{\widehat{p} | p \in P, p \cdot x = y\}.$$

This category is isomorphic to $\mathcal{B}_P X$; the claim is that the inclusion $\check{\mathcal{B}}(P) \subseteq \mathcal{G}(P)$ induces a deformation retract after realization. This will follow if we can find a functor $\Psi : \mathcal{G}(P) \rightarrow \check{\mathcal{B}}(P)$ that is the identity on $\check{\mathcal{B}}(P)$ together with a natural transformation $F : \text{Id}_{\mathcal{G}(P)} \Rightarrow \iota_{\check{\mathcal{B}}(P)}^{\mathcal{G}(P)} \circ \Psi$, where $\iota_{\check{\mathcal{B}}(P)}^{\mathcal{G}(P)} : \check{\mathcal{B}}(P) \rightarrow \mathcal{G}(P)$ is the inclusion functor.

Pick some lifting of morphisms $\xi : \text{Mor}(\mathcal{O}^{c\mathfrak{X}}) \rightarrow \text{Mor}(\mathcal{L}^{\mathfrak{X}})$, and assume that ξ sends identities to identities. Let $\Psi : \mathcal{G}(P) \rightarrow \check{\mathcal{B}}(P)$ be the functor

$$\begin{array}{ccc} ((Q, \alpha), x) & \dashrightarrow & ((P, \text{id}_P), \ell_{\xi(\alpha)}(x)) \\ \mathbf{g} \downarrow & & \downarrow \Psi(\mathbf{g}) \\ ((R, \beta), y) & \dashrightarrow & ((P, \text{id}_P), \ell_{\xi(\beta)}(y)) \end{array}$$

where $\Psi(\mathbf{g})$ is defined to be \widehat{p} for the unique $p \in P$ such that

$$\begin{array}{ccc} Q & \xrightarrow{\mathbf{g}} & R \\ \xi(\alpha) \downarrow & & \downarrow \xi(\beta) \\ P & \xrightarrow{\widehat{p}} & P \end{array}$$

The existence and uniqueness of p follow from Proposition 3.2.18 and the fact that $\overline{\pi}(\xi(\beta) \circ \mathbf{g}) = \beta \circ \overline{\pi}(\mathbf{g}) = \alpha = \overline{\pi}(\xi(\alpha))$ because \mathbf{g} is a morphism in $(\overline{\pi} \downarrow P)$.

We must check that \widehat{p} indeed defines a morphism in $\mathcal{G}(P)$:

$$((P, \text{id}_P), \ell_{\xi(\alpha)}(x)) \rightarrow ((P, \text{id}_P), \ell_{\xi(\beta)}(y)).$$

To do this, we must show that

$$\begin{array}{ccc} P & \xrightarrow{(c_p, \ell_p)} & P \\ & \searrow \text{id}_P & \swarrow \text{id}_P \\ & P & \end{array}$$

commutes in $\mathcal{O}^{c\mathfrak{X}}$ and that $\ell_p \circ \ell_{\xi(\alpha)}(x) = \ell_{\xi(\beta)}(y)$. The first is obvious from the definition of $\mathcal{O}^{c\mathfrak{X}}$, and the second follows from the fact that $\widehat{p} \circ \xi(\alpha) = \xi(\beta) \circ \mathbf{g}$, so that $\ell_p \circ \ell_{\xi(\alpha)}(x) = \ell_{\xi(\beta)} \circ \ell_{\mathbf{g}}(x) = \ell_{\xi(\beta)}(y)$ by the assumption on \mathbf{g} .

Observe that $\Psi|_{\check{\mathcal{B}}(P)}$ is the identity functor.

Define the natural transformation $\Theta : \text{Id}_{\mathcal{G}(P)} \Rightarrow \iota_{\check{\mathcal{B}}(P)}^{\mathcal{G}(P)} \circ \Psi$ by

$$\Theta(((Q, \alpha), x)) = \xi(\alpha) : ((Q, \alpha), x) \rightarrow ((P, \text{id}_P), \ell_{\xi(\alpha)}(x)).$$

That this is a morphism in $\mathcal{G}(P)$ follows easily from the definition of $\mathcal{G}(P)$ above and the fact that $\xi(\alpha)$ is a lift of α to \mathcal{L} . If $\mathfrak{g} : ((Q, \alpha), x) \rightarrow ((R, \beta), y)$ is a morphism in $\mathcal{G}(P)$, we want the diagram

$$\begin{array}{ccc} ((Q, \alpha), x) & \xrightarrow{\mathfrak{g}} & ((R, \beta), y) \\ \xi(\alpha) \downarrow & & \downarrow \xi(\beta) \\ ((P, \text{id}_P), \ell_{\xi(\alpha)}(x)) & \xrightarrow{\Psi(\mathfrak{g})} & ((P, \text{id}_P), \ell_{\xi(\beta)}(y)) \end{array}$$

to commute in $\mathcal{L}^{\mathfrak{X}}$, which follows from the definition of $\Psi(\mathfrak{g})$.

Finally, note that for any object of $\tilde{\mathcal{B}}(P)$, it is easy to see that $F(\tilde{x})$ is the identity. Thus the geometric realization of Θ gives a homotopy that shows that $|\Psi|$ realizes $|\tilde{\mathcal{B}}(P)|$ as a deformation retract of $|\mathcal{G}(P)|$, and we have the result on objects.

On morphisms: First observe that each $[\varphi, \sigma] \in \mathcal{O}^{c\mathfrak{X}}(P, P')$ induces a functor $(\bar{\pi} \downarrow P) \rightarrow (\bar{\pi} \downarrow P')$:

$$\begin{array}{ccc} (Q, \alpha) & \dashrightarrow & (Q, [\varphi, \sigma] \circ \alpha) \\ \mathfrak{g} \downarrow & & \downarrow \mathfrak{g} \\ (R, \beta) & \dashrightarrow & (R, [\varphi, \sigma] \circ \beta) \end{array}$$

This in turn induces a functor $[\varphi, f]_* : \mathcal{G}(P) \rightarrow \mathcal{G}(P')$, defined by

$$\begin{array}{ccc} ((Q, \alpha), x) & \dashrightarrow & ((Q, [\varphi, \sigma] \circ \alpha), x) \\ \mathfrak{g} \downarrow & & \downarrow \mathfrak{g} \\ ((R, \beta), \ell_{\mathfrak{g}} x) & \dashrightarrow & ((R, [\varphi, \sigma] \circ \beta), \ell_{\mathfrak{g}} x) \end{array}$$

Similarly, for $(\varphi, \sigma) \in \mathfrak{X}(P, P')$ a lift of $[\varphi, \sigma] \in \mathcal{O}^{c\mathfrak{X}}(P, P')$, there is a functor $\tilde{\mathcal{B}}_{(\varphi, \sigma)} : \tilde{\mathcal{B}}(P) \rightarrow \tilde{\mathcal{B}}(P')$:

$$\begin{array}{ccc} ((P, \text{id}_P), x) & \dashrightarrow & ((P', \text{id}_{P'}), \sigma(x)) \\ \widehat{p} \downarrow & & \downarrow \widehat{\varphi p} \\ ((P, \text{id}_P), p \cdot x) & \dashrightarrow & ((P', \text{id}_{P'}), \varphi p \cdot \sigma(x)) \end{array}$$

To finish the proof that \tilde{B} is a homotopy lifting of B we must show that the diagram of functors

$$\begin{array}{ccc} \tilde{\mathcal{B}}(P) & \xrightarrow{\iota} & \mathcal{G}(P) \\ \tilde{\mathcal{B}}_{(\varphi, \sigma)} \downarrow & & \downarrow [\varphi, \sigma]_* \\ \tilde{\mathcal{B}}(P') & \xrightarrow{\iota} & \mathcal{G}(P') \end{array}$$

commutes up to natural transformation.

Let F_1 be the functor given by the top path of the diagram. Explicitly, F_1 is the functor

$$\begin{array}{ccc} ((P, \text{id}_P), x) & \dashrightarrow & ((P, [\varphi, \sigma]), x) \\ \widehat{p} \downarrow & & \downarrow \widehat{p} \\ ((P, \text{id}_P), p \cdot x) & \dashrightarrow & ((P, [\varphi, \sigma]), p \cdot x) \end{array}$$

Let the bottom path be the functor F_2 :

$$\begin{array}{ccc} ((P, \text{id}_P), x) & \dashrightarrow & ((P', \text{id}_{P'}), \sigma(x)) \\ \hat{p} \downarrow & & \downarrow \widehat{\varphi p} \\ ((P, \text{id}_P), p \cdot x) & \dashrightarrow & ((P', \text{id}_{P'}), \varphi p \cdot \sigma(x)) \end{array}$$

Choose some $\mathfrak{h} \in \mathcal{L}^{\mathfrak{X}}(P, P')$ that lifts (φ, σ) , and define the natural transformation

$$\Phi : F_1 \Rightarrow F_2 : ((P, \text{id}), x) \mapsto \mathfrak{h}.$$

To see that \mathfrak{h} is indeed a morphism in $\mathcal{G}(P')$ from $((P, [\varphi, \sigma]), x)$ to $((P', \text{id}_{P'}), \sigma(x))$, note that

$$\begin{array}{ccc} P & \xrightarrow{\pi(\mathfrak{h})} & P' \\ & \searrow [\varphi, f] & \swarrow \text{id}_{P'} \\ & P' & \end{array}$$

commutes because \mathfrak{h} is a lifting of $[\varphi, \sigma]$. Moreover, $\ell_{\mathfrak{h}} x = \sigma(x)$ because \mathfrak{h} is a lifting of (φ, σ) .

Finally, if \hat{p} is a morphism in $\tilde{\mathcal{B}}(P)$ from $((P, \text{id}_P), x)$ to $((P, \text{id}_P), y)$, the diagram

$$\begin{array}{ccc} ((P, [\varphi, \sigma]), x) & \xrightarrow{\mathfrak{h}} & ((P', \text{id}_{P'}), \sigma(x)) \\ \hat{p} \downarrow & & \downarrow \widehat{\varphi p} \\ ((P, [\varphi, \sigma]), y) & \xrightarrow{\mathfrak{h}} & ((P', \text{id}_{P'}), \sigma(y)) \end{array}$$

commutes in $\mathcal{G}(P')$ because Axiom (C) and the fact that $c_{\mathfrak{h}} = \varphi$ (\mathfrak{h} lifts (φ, σ)) together imply that $\widehat{\varphi p} \circ \mathfrak{h} = \mathfrak{h} \circ \hat{p}$ in $\mathcal{L}^{\mathfrak{X}}(P, P')$.

This completes the proof that \tilde{B} is a homotopy lifting of B . \square

3.3 Obstruction theory

Fix an abstract fusion action system \mathfrak{X} . The material in this section does not depend on \mathfrak{X} 's being saturated.

We can solve the homotopy lifting problem for $\tilde{\mathcal{B}}_- X : \mathcal{O}^{c\mathfrak{X}} \rightarrow \text{ho}\mathcal{TOP}$, and thus create a space that we might call a “classifying space” for \mathfrak{X} , if we have an associated linking action system $\mathcal{L}^{\mathfrak{X}}$. In this section we describe how we can construct $\mathcal{L}^{\mathfrak{X}}$ from the data of \mathfrak{X} , or more accurately, we describe the difficulties in doing so. We also describe the obstructions to constructing $\mathcal{L}^{\mathfrak{X}}$ uniquely.

Definition 3.3.1. If there is a unique $\mathcal{L}^{\mathfrak{X}}$ associated to \mathfrak{X} , the space $\text{hocolim}_{\mathcal{L}^{\mathfrak{X}}} \mathbb{X}$ is the *classifying space* for the fusion action system. We denote this space by $B\mathfrak{X}$.

Notation 3.3.2. We denote by $(\tilde{\varphi}, \tilde{\sigma}) \in \mathfrak{X}(P, Q)$ a lift of $[\varphi, \sigma] \in \mathcal{O}^{c\mathfrak{X}}(P, Q)$.

Definition 3.3.3. Let $\mathcal{Z}^{\mathfrak{X}} : (\mathcal{O}^{c\mathfrak{X}})^{op} \rightarrow \mathcal{Ab}$ be the functor

$$\begin{array}{ccccc} P & \dashrightarrow & Z(P; X) & \xleftarrow{\tilde{\varphi}^{-1}} & Z(\tilde{\varphi}(P); X) = Z_S(\tilde{\varphi}(P); X) \\ \downarrow [\varphi, \sigma] & & \uparrow \mathcal{Z}^{\mathfrak{X}}([\varphi, \sigma]) & & \uparrow \text{incl} \\ Q & \dashrightarrow & Z(Q; X) & \xlongequal{\quad\quad\quad} & Z_S(Q; X) \end{array}$$

where $\mathcal{Z}^{\mathfrak{X}}([\varphi, \sigma])$ is the unique map of groups that makes the rectangle commute.

Proposition 3.3.4. *The functor $Z^{\mathfrak{X}}$ is well-defined.*

Proof. The equalities come from the assumption that P is X -centric. We must simply show that for $(\tilde{\varphi}', \tilde{\sigma}')$ any another lift of $[\varphi, \sigma]$, the maps $\tilde{\varphi}^{-1}$ and $(\tilde{\varphi}')^{-1}$ are equal on $Z(Q; X)$. Since that $\tilde{\varphi}$ and $\tilde{\varphi}'$ are both lifts of φ , there is some $q \in Q$ such that $\tilde{\varphi}' = c_q \circ \tilde{\varphi}$. Therefore c_q^{-1} is the identity on $Z(Q; X)$, and the result follows. \square

Just as in the theory of p -local finite groups, there is a well-defined obstruction theory that governs the existence and uniqueness of linking action systems associated to a fusion action system.

Theorem 3.3.5. *The data of the abstract fusion action \mathfrak{X} determines an element $u \in \lim_{\mathcal{O}^{c\mathfrak{X}}}^3 Z^{\mathfrak{X}}$ that vanishes precisely when there is a linking action system $\mathcal{L}^{\mathfrak{X}}$ associated to \mathfrak{X} . Moreover, if there is a linking action system, the group $\lim_{\mathcal{O}^x}^2 Z^{\mathfrak{X}}$ acts transitively on the set of linking action systems viewed as categories over $\mathcal{O}^{c\mathfrak{X}}$.*

Proof. Just as the proof of Theorem 2.6.15 looks almost exactly like that of the analogue for fusion systems, so too can the proof of the obstruction theory governing p -local finite group from [BLO2] be carried over with only the obvious changes to the world of p -local finite group actions. Again, we omit reproducing what amounts to identical symbols with a different interpretation, and instead refer the reader to the original source. \square

3.4 Linking action systems as transporter systems

Let \mathfrak{X} be a saturated fusion action system and $\mathcal{L}^{\mathfrak{X}}$ an associated linking action system.

Proposition 3.2.6 can be thought of a “unique right lifting” lemma, and we saw that it implied many useful properties for a linking action system $\mathcal{L}^{\mathfrak{X}}$. In particular, we learned from it that all morphisms are categorically mono, that restriction is a well-defined notion (given the existence of specified “inclusion morphisms”), and ultimately that all morphisms are epi. This last result “should” have been derived from a “unique left lifting” lemma, but instead we used the additional structure of the inclusion morphisms to derive it. Indeed, there is no direct left lifting analogue of Proposition 3.2.6; instead, we have to settle with the following, which turns out to have its own uses:

Proposition 3.4.1. *Let $P \xrightarrow{(\varphi, \sigma)} Q \xrightarrow{(\psi, \tau)} R$ be a sequence of morphisms in \mathfrak{X} . For any*

$$\mathfrak{h} \in \pi_{P, Q}^{-1}((\varphi, \sigma)) \subseteq \mathcal{L}^{\mathfrak{X}}(P, Q) \quad \text{and} \quad \widetilde{\mathfrak{gh}} \in \pi_{P, R}^{-1}((\psi\varphi, \tau\sigma)) \subseteq \mathcal{L}^{\mathfrak{X}}(P, R)$$

there is a unique $\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(P, Q)$ such that $\mathfrak{gh} = \widetilde{\mathfrak{gh}}$. Moreover, there is a unique $z \in \varphi(Z(P; X))$ such that $(c_{\mathfrak{g}}, \ell_{\mathfrak{g}}) = (\varphi c_z, \sigma \ell_z)$.

Proof. The morphism \mathfrak{h} is epi by Proposition 3.2.14, so there is at most one \mathfrak{g} such that $\mathfrak{gh} = \widetilde{\mathfrak{gh}}$.

Pick any \mathfrak{g}' lifting (ψ, τ) . Thus $\mathfrak{g}' \circ \mathfrak{h}$ and $\widetilde{\mathfrak{gh}}$ have the same image in $\mathfrak{X}(P, R)$, and by Axioms (A) and (C), there is a unique $z \in Z(P; X)$ such that $\widetilde{\mathfrak{gh}} = \mathfrak{g}' \circ \mathfrak{h} \circ \widehat{z} = \mathfrak{g}' \circ \widehat{\varphi(z)} \circ \mathfrak{h}$. Then $\mathfrak{g} = \mathfrak{g}' \circ \widehat{\varphi(z)}$ gives us the existence statement. \square

Remark 3.4.2. The main difference between the left “lifting” of Proposition 3.4.1 and the right lifting of Proposition 3.2.6 is that this more recent result cannot actually lift the morphism $(\psi, \tau) \in \mathfrak{X}(Q, R)$, but only some $\varphi(Z(P; X))$ -translate of it. The difference between the two stems from the possibility that $Q \not\geq \varphi(P)$, in which case φ need not take central elements to central elements; as Q can be bigger, it is possible that $\varphi(z)$ acts nontrivially on Q , even though z acts trivially on P .

Recall that an *extension* of $\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(P, Q)$ is $\tilde{\mathfrak{g}} \in \mathcal{L}^{\mathfrak{X}}(\tilde{P}, \tilde{Q})$ for $P \leq \tilde{P}$ and $Q \leq \tilde{Q}$ such that

$$\begin{array}{ccc} \tilde{P} & \xrightarrow{\tilde{\mathfrak{g}}} & \tilde{Q} \\ \uparrow i_{\tilde{P}}^{\tilde{P}} & & \uparrow i_{\tilde{Q}}^{\tilde{Q}} \\ P & \xrightarrow{\mathfrak{g}} & Q \end{array}$$

commutes in $\mathcal{L}^{\mathfrak{X}}$. We have already seen (thanks to the fact that all morphisms of $\mathcal{L}^{\mathfrak{X}}$ are both epi and mono) that extensions are unique if they exist, and we are now in the position to say when exactly they do exist:

Proposition 3.4.3. *Let $\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(P, Q)$ be an isomorphism and let $\tilde{P}, \tilde{Q} \leq S$ be such that $P \trianglelefteq \tilde{P}$, $Q \trianglelefteq \tilde{Q}$, and $\mathfrak{g} \circ \delta_{P,P}(\tilde{P}) \circ \mathfrak{g}^{-1} \leq \delta_{Q,Q}(\tilde{Q})$. Then there is a unique extension $\tilde{\mathfrak{g}} \in \mathcal{L}^{\mathfrak{X}}(\tilde{P}, \tilde{Q})$ of \mathfrak{g} .*

Proof. First suppose that Q is fully X -centralized. By the Extension Axiom for saturated fusion actions, the morphism $(c_{\mathfrak{g}}, \ell_{\mathfrak{g}}) \in \mathfrak{X}(P, Q)$ extends to some $(\varphi, \ell_{\mathfrak{g}}) \in \mathfrak{X}(N_{(c_{\mathfrak{g}}, \ell_{\mathfrak{g}})}, S)$. The condition on \tilde{P} implies that $\tilde{P} \leq N_{(c_{\mathfrak{g}}, \ell_{\mathfrak{g}})}$ (project the condition down to \mathfrak{X}) and the condition on \tilde{Q} implies that $\varphi(p') \in \tilde{Q}$ for all $p' \in \tilde{P}$. We

can thus rename $(\varphi, \ell_{\mathfrak{g}})$ to be its restriction in $\mathfrak{X}(\tilde{P}, \tilde{Q})$. For the sequence $P \xrightarrow{(\iota_{\tilde{P}}^{\tilde{P}}, \text{id}_X)} \tilde{P} \xrightarrow{(\varphi, \ell_{\mathfrak{g}})} \tilde{Q}$ in \mathfrak{X} , let $i_{\tilde{P}}^{\tilde{P}} \in \mathcal{L}^{\mathfrak{X}}(P, \tilde{P})$ lift the first map and $i_{\tilde{Q}}^{\tilde{Q}} \circ \mathfrak{g}$ lift the composite. Then Proposition 3.4.1 applies to give a unique extension $\tilde{\mathfrak{g}}$ of \mathfrak{g} as desired (though note that it need not be the case that $\tilde{\mathfrak{g}}$ is a lift of $(\varphi, \ell_{\mathfrak{g}})$, only a lift of a $\varphi(Z(P; X))$ -translate of it).

Now consider the general case, where Q need not be fully X -centralized. Let R be fully normalized and \mathcal{F} -conjugate to P and Q . For any $\mathfrak{h} \in \mathcal{L}^{\mathfrak{X}}(Q, R)$, we have that $\mathfrak{h} \circ \widehat{N_S(Q)}|_Q^Q \circ \mathfrak{h}^{-1} \leq \mathcal{L}^{\mathfrak{X}}(R)$ is an inclusion of a p -subgroup. As R is fully normalized, Proposition 3.2.16 states that $\delta_{R,R}(N_S(R))$ is Sylow in $\mathcal{L}^{\mathfrak{X}}(R)$, so \mathfrak{h} can be chosen so that $\mathfrak{h} \circ \widehat{N_S(Q)}|_Q^Q \circ \mathfrak{h}^{-1} \leq \widehat{N_S(R)}|_R^R$. The subgroup R is fully X -centralized, so the first part of this proof implies that there are morphisms $\tilde{\mathfrak{h}}$ extending \mathfrak{h} to $N_S(Q)$ and $\tilde{\mathfrak{h}}\mathfrak{g}$ extending $\mathfrak{h}\mathfrak{g}$ to \tilde{P} . Let $\bar{\mathfrak{h}}$ be the restricted isomorphism of $\tilde{\mathfrak{h}}$ with source \tilde{Q} , and similarly $\bar{\mathfrak{h}}\mathfrak{g}$ the restricted isomorphism of $\tilde{\mathfrak{h}}\mathfrak{g}$ with source \tilde{P} . The situation can be represented as:

$$\begin{array}{ccc} N_S(Q) & \xrightarrow{\tilde{\mathfrak{h}}} & N_S(R) \\ \uparrow i & & \uparrow i \\ \tilde{Q} & \xrightarrow[\bar{\mathfrak{h}}]{\cong} & c_{\tilde{\mathfrak{h}}}(\tilde{Q}) \\ \uparrow i & & \uparrow i \\ Q & \xrightarrow[\mathfrak{h}]{\cong} & R \end{array} \quad \begin{array}{ccc} \tilde{P} & \xrightarrow{\tilde{\mathfrak{h}}\mathfrak{g}} & N_S(R) \\ \uparrow i & & \uparrow i \\ \tilde{P} & \xrightarrow[\bar{\mathfrak{h}}\mathfrak{g}]{\cong} & c_{\tilde{\mathfrak{h}}\mathfrak{g}}(\tilde{P}) \\ \uparrow i & & \uparrow i \\ P & \xrightarrow[\mathfrak{h} \circ \mathfrak{g}]{} & R \end{array}$$

The claim is that $c_{\bar{\mathfrak{h}}\mathfrak{g}}(\tilde{P}) \leq c_{\bar{\mathfrak{h}}}(\tilde{Q})$. Axiom (C) implies that for $p' \in \tilde{P}$ and $q' \in \tilde{Q}$,

$$\widehat{c_{\bar{\mathfrak{h}}\mathfrak{g}}(p')} = \bar{\mathfrak{h}}\mathfrak{g} \circ \hat{p}' \circ \bar{\mathfrak{h}}\mathfrak{g}^{-1} \quad \text{and} \quad \widehat{c_{\bar{\mathfrak{h}}}(\tilde{Q})} = \bar{\mathfrak{h}} \circ \hat{q}' \circ \bar{\mathfrak{h}}^{-1}.$$

Observing that every $p' \in \tilde{P}$ defines a morphism in $\mathcal{L}^{\mathfrak{X}}(P)$, we can restrict this to get

$$\delta_{R,R}(c_{\bar{\mathfrak{h}}\mathfrak{g}}(\tilde{P})) = (\mathfrak{h} \circ \mathfrak{g}) \circ \delta_{P,P}(\tilde{P}) \circ (\mathfrak{h} \circ \mathfrak{g})^{-1} \leq \mathfrak{h} \circ \delta_{Q,Q}(\tilde{Q}) \circ \mathfrak{h}^{-1} = \delta_{R,R}(c_{\bar{\mathfrak{h}}}(\tilde{Q}))$$

where the inequality comes from the initial assumption on \tilde{P} and \tilde{Q} , and the claim is proved.

Therefore the Divisibility Axiom of fusion action systems implies that there is some $(\psi, \tau) \in \mathfrak{X}(\tilde{P}, \tilde{Q})$ such that $c_{\tilde{\mathfrak{h}}\tilde{\mathfrak{g}}} = (\iota, \text{id}_X) \circ c_{\tilde{\mathfrak{h}}} \circ (\psi, \tau) \in \mathfrak{X}(\tilde{P}, N_S(R))$. Now Proposition 3.2.6 implies that there is a unique $\tilde{\mathfrak{g}} \in \mathcal{L}^{\mathfrak{X}}(\tilde{P}, \tilde{Q})$ such that $\tilde{\mathfrak{h}}\tilde{\mathfrak{g}} = \tilde{\mathfrak{h}}\tilde{\mathfrak{g}}$. Restricting this to P we get $\mathfrak{h} \circ \mathfrak{g} = \mathfrak{h} \circ \text{res}_P^Q(\tilde{\mathfrak{g}})$, and the fact that \mathfrak{h} is mono implies that $\text{res}_P^Q(\tilde{\mathfrak{g}}) = \mathfrak{g}$, as desired. \square

Remark 3.4.4. Note that the condition on the overgroups can be restated as follows: For every $p' \in \tilde{P}$, there is some $q' \in \tilde{Q}$ such that $\mathfrak{g} \circ p' \circ \mathfrak{g}^{-1} = q'$. This condition is morally the same as the definition of the extender $N_{(\varphi, \sigma)}$ used to state the Extension Axiom for saturated fusion actions. The key difference is that, when stated in terms of the linking action system, the extension condition on the source is sharper, which allows us to relax the assumption that the target of \mathfrak{g} be fully X -centralized.

Recall the notion of an *abstract transporter system* associated to a fusion system \mathcal{F} , as introduced in [OV] and briefly described in Subsection 1.4.

Corollary 3.4.5. *For a linking action system $\mathcal{L}^{\mathfrak{X}}$ associated to \mathfrak{X} the composite $\mathcal{L}^{\mathfrak{X}} \rightarrow \mathfrak{X} \rightarrow \mathcal{F}$, together with $\delta : \mathcal{T}_S^{c_X} \rightarrow \mathcal{L}^{\mathfrak{X}}$, give $\mathcal{L}^{\mathfrak{X}}$ the structure of a transporter system associated to \mathcal{F} .*

Proof. The only difficult part of the proof is the extension condition, which Proposition 3.4.3 shows to be true. \square

We can interpret this result as saying that a fusion action system \mathfrak{X} and an associated linking system $\mathcal{L}^{\mathfrak{X}}$ give rise to a transporter system on the underlying fusion system \mathcal{F} together with a map $\text{Mor}(\mathcal{L}^{\mathfrak{X}}) \rightarrow \Sigma_X$ that takes composition to multiplication and inclusions to the identity. In [OV] it is shown that this map is equivalent to the data of a group map $\pi_1(|\mathcal{L}^{\mathfrak{X}}|) \rightarrow \Sigma_X$. We now set out to reverse this process:

Fix a (saturated) fusion system \mathcal{F} on the p -group S and a transporter system \mathcal{T} associated to \mathcal{F} . Assume that we are given a group map $\theta : \pi_1(|\mathcal{T}|) \rightarrow \Sigma_X$, or equivalently a map $\text{Mor}(\mathcal{T}) \rightarrow \Sigma_X$ that sends composition to multiplication and inclusions to the identity.

Definition 3.4.6. Let \mathfrak{X}^θ be the category with $\text{Ob}(\mathfrak{X}^\theta) = \text{Ob}(\mathcal{T})$ and morphisms given by

$$\mathfrak{X}^\theta(P, Q) = \{(\varphi, \sigma) \in \text{Inj}(P, Q) \times \Sigma_X \mid \exists \mathfrak{g} \in \mathcal{T}(P, Q) \text{ such that } (\varphi, \sigma) = (c_{\mathfrak{g}}, \theta(\mathfrak{g}))\}.$$

This allows us to define an action of S on X as follows: The p -group S embeds as a subgroup of $\mathcal{T}(S)$ via the structure map $\delta : \mathcal{T}_S^{\text{Ob}(\mathcal{T})} \rightarrow \mathcal{T}$. We denote by \hat{S} the image of S . Moreover, θ defines a $\mathcal{T}(S)$ -action on X , and thus an S -action by restriction.

Clearly \mathfrak{X}^θ is a fusion action system, or at least generates one once we allow for restrictions of morphisms to subgroups not in $\text{Ob}(\mathcal{T})$. Let \mathcal{F}^θ be the underlying fusion system.

We would for \mathfrak{X}^θ to be saturated, but for now we must settle for a weaker condition.

Definition 3.4.7. For \mathcal{C} a collection of subgroups of S closed under \mathcal{F}^θ -conjugacy and overgroups, and \mathfrak{X} a fusion action system on S , we say that \mathfrak{X} is $\text{Ob}(\mathcal{C})$ -saturated if the Saturation Axioms hold for all $P \in \mathcal{C}$.

We need a little terminology to prove object-saturation of \mathcal{F}^θ :

Notation 3.4.8. In the above situation, for any $P \in \text{Ob}(\mathcal{T})$ we define $E(P) = \ker[\mathcal{T}(P) \rightarrow \mathcal{F}^\theta(P)]$. We also denote by $K(P)$ the kernel of the action map $\theta : \mathcal{T}(P) \rightarrow \Sigma(X)$. Finally, let C be the core of the S -action on X , so that $\hat{C} = \hat{S} \cap K(S)$. We can therefore define the notions of X -normalizers and X -centralizers of objects of \mathcal{T} in the obvious way.

Proposition 3.4.9. *For each $P \in \text{Ob}(\mathcal{T})$,*

- *P is fully normalized in \mathcal{F}^θ if and only if $\widehat{N_S(P)} \in \text{Syl}_p(\mathcal{T}(P))$.*
- *P is fully centralized in \mathcal{F}^θ if and only if $\widehat{Z_S(P)} \in \text{Syl}_p(E(P))$.*

- P is fully X -normalized in \mathcal{F}^θ if and only if $N_S(\widehat{P}; X) \in \text{Syl}_p(K(P))$.
- P is fully X -centralized in \mathcal{F}^θ if and only if $Z_S(\widehat{P}; X) \in \text{Syl}_p(E(P) \cap K(P))$.

Proof. The first two points are proved in [OV], Proposition 3.4. The proofs of the remaining two points follow basically the same argument as the second.

Proof of third point: For $P \in \text{Ob}(\mathcal{T})$, let Q be \mathcal{F}^θ -conjugate to P and fully normalized, so by the first point $N_S(\widehat{Q}) \in \text{Syl}_p(\mathcal{T}(Q))$. Therefore

$$N_S(\widehat{Q}; X) = N_S(\widehat{Q}) \cap K(Q) \in \text{Syl}_p \mathcal{T}(Q).$$

Now, \mathcal{F}^θ is $\text{Ob}(\mathcal{T})$ -saturated by [OV], so the proof of Proposition ?? applies here to give us that P is also fully X -normalized. Since $K(Q) \cong K(P)$, we have $N_S(\widehat{P}; X) \in \text{Syl}_p(K(P))$ if and only if $|N_S(P; X)| = |N_S(Q; X)|$, or equivalently, if and only if P is fully X -normalized.

The proof of the fourth point is the same as that of the third, replacing every instance of $K(-)$ with $E(-) \cap K(-)$. \square

The notation in the following Corollary is a direct analogy with that introduced to describe fusion action systems.

Corollary 3.4.10. *For all $P \in \text{Ob}(\mathcal{T})$*

- *If P is fully normalized then*
 - $\mathfrak{X}_S^\theta(P) \in \text{Syl}_p(\mathfrak{X}^\theta(P))$.
 - $\mathcal{F}_S^\theta(P) \in \text{Syl}_p(\mathcal{F}^\theta(P))$.
 - $\Sigma_\theta^S(P) \in \text{Syl}_p(\Sigma_\theta(P))$.
- *If P is fully X -normalized, then $\mathcal{F}_S^\theta(P)_0 \in \text{Syl}_p(\mathcal{F}^\theta(P)_0)$.*
- *If P is fully centralized, then $\Sigma_\theta^S(P)_0 \in \text{Syl}_p(\Sigma_\theta(P)_0)$.*

Proof. Each of these follows from the observation that the image of a Sylow is Sylow in the quotient. \square

Proposition 3.4.11. *Let $Q \in \text{Ob}(\mathcal{T})$ be fully X -centralized and $(\varphi, \sigma) \in \text{Iso}_{\mathfrak{X}^\theta}(P, Q)$. Then there is some $(\tilde{\varphi}, \sigma) \in \mathfrak{X}^\theta(N_{(\varphi, \sigma)}, S)$ that extends (φ, σ) .*

Proof. We first claim that $N_S(\widehat{Q}) \in \text{Syl}_p(\widehat{N_S(Q)} \cdot E(Q) \cap K(Q))$:

$$\begin{aligned} \left[\widehat{N_S(Q)} : \widehat{N_S(Q)} \cdot E(Q) \cap K(Q) \right] &= \frac{|\widehat{N_S(Q)}| \cdot |E(Q) \cap K(Q)|}{|N_S(Q)| \cdot |\widehat{N_S(Q)} \cap E(Q) \cap K(Q)|} \\ &= \left[\widehat{N_S(Q)} \cap E(Q) \cap K(Q) : E(Q) \cap K(Q) \right]. \end{aligned}$$

The fact that $\widehat{N_S(Q)} \cap E(Q) \cap K(Q) = Z_S(Q; X)$ and the final point of Proposition 3.4.9 gives the claim.

Now, pick $\mathfrak{g} \in \mathcal{T}(P, Q)$ such that $(c_{\mathfrak{g}}, \theta(\mathfrak{g})) = (\varphi, \sigma)$. By definition of $N_{(\varphi, \sigma)}$,

$$\mathfrak{g} \circ \widehat{N_{(\varphi, \sigma)}}|_P^P \circ \mathfrak{g}^{-1} \leq \widehat{N_S(Q)}|_Q^Q \cdot E(Q) \cap K(Q)$$

and so by the Sylow result just proved, there is some $\mathfrak{h} \in E(Q) \cap K(Q)$ such that

$$(\mathfrak{h}\mathfrak{g}) \circ \widehat{N_{(\varphi, \sigma)}}|_P^P \circ (\mathfrak{h}\mathfrak{g})^{-1} \leq \widehat{N_S(Q)}|_Q^Q.$$

Then $\mathfrak{h}\mathfrak{g} \in \mathcal{T}(P, Q)_{\text{iso}}$, $N_{(\varphi, \sigma)} \supseteq P$, and $N_S(Q) \triangleright Q$, so the conditions for the Extension Axiom (II) of abstract transporter systems are satisfied. Therefore there is some $\tilde{\mathfrak{h}\mathfrak{g}} \in \mathcal{T}(N_{(\varphi, \sigma)}, N_S(Q))$ that extends $\mathfrak{h}\mathfrak{g}$. This implies that $(c_{\mathfrak{h}\mathfrak{g}}, \theta(\mathfrak{h}\mathfrak{g})) \in \mathfrak{X}^\theta(N_{(\varphi, \sigma)}, S)$ extends $(c_{\mathfrak{g}}, \theta(\mathfrak{g})) = (\varphi, \sigma)$ in \mathfrak{X}^θ , and the result is proved. \square

Corollary 3.4.12. *The fusion action system \mathfrak{X}^θ is $\text{Ob}(\mathcal{T})$ -saturated.*

Proof. All the Axioms have been verified in Propositions 3.4.9 and 3.4.11. \square

Remark 3.4.13. Corollary 3.4.12 suggests that we should consider under what circumstances a collection of subgroups \mathcal{H} of S has the property that if a fusion action system \mathfrak{X} is \mathcal{H} -saturated, then \mathfrak{X} is itself saturated. Such a result for fusion systems can be found in [BCG⁺], suggesting that it is not unreasonable to expect that the same would apply for fusion action systems. Hopefully this point will be addressed in future research.

We have seen that a transporter system \mathcal{T} together with a map $\theta : \pi_1(\mathcal{T}) \rightarrow \Sigma_X$ determine a saturated fusion action system, at least so far as the objects of \mathcal{T} are aware. We are given natural functors $\mathcal{T}_S^{\text{Ob}(\mathcal{T})} \longrightarrow \mathcal{T} \longrightarrow \mathfrak{X}^\theta$, and we can ask how close this is to being the data of an X -centric linking action system associated to \mathfrak{X}^θ .

The following result states that, so long as all the X -centric subgroups are accounted for, \mathcal{T} fails to be a linking action system “in a p' -way,” and moreover that it contains enough data to construct a linking action system \mathcal{L}^θ :

Proposition 3.4.14. *Suppose that in the above situation $\text{Ob}(\mathcal{T})$ contains all X -centric subgroups of S , and let \mathcal{T}^{cX} denote the fully subcategory with these as the objects. Then for any X -centric $P \leq S$, there is a unique p' -group $EK'(P)$ such that $E(P) \cap K(P) = \widehat{Z(P; X)} \times EK'(P)$. Furthermore, $EK'(P)$ is the subgroup of all p' -elements of $E(P) \cap K(P)$.*

Consequently, if \mathcal{L}^θ is the category whose objects are the X -centric subgroups of S and whose morphisms are given by

$$\mathcal{L}^\theta(P, Q) = \mathcal{T}(P, Q) / EK'(P),$$

then \mathcal{L}^θ is an X -centric linking action system associated to \mathfrak{X}^θ .

Proof. Axiom (C) of transporter systems implies that $E(P)$ commutes with \widehat{P} , so in particular $E(P) \cap K(P)$ does as well. Therefore the fact that P is X -centric implies that

$$\widehat{S} \cap E(P) \cap K(P) = \widehat{Z_S(P; X)} = \widehat{Z(P; X)} \trianglelefteq E(P) \cap K(P),$$

and the fact that P is fully X -centralized implies that $\widehat{Z(P; X)}$ is a normal abelian Sylow subgroup of $E(P) \cap K(P)$. The Schur-Zassenhaus theorem then implies the existence and uniqueness of $EK'(P)$, from which it easily follows that \mathcal{L}^θ is an X -centric linking action system associated to \mathfrak{X}^θ . \square

3.5 Stabilizers of p -local finite group actions

For $H \leq G$ finite groups, it is a basic result that

$$BH \simeq EG \times_H * \simeq EG \times_G G/H.$$

In this section we prove that the analogous statement for linking actions systems is true. Let \mathfrak{X} be a saturated fusion action system, and let $\mathcal{L}^\mathfrak{X}$ be an associated linking action system. Recall that C denotes the core of the S action on X .

Recall that the fusion action system \mathfrak{X} is *transitive* if $\mathfrak{X}(1)$ acts transitively on X .

Lemma 3.5.1. *If the fusion action system \mathfrak{X} is saturated, $\pi_\Sigma(\mathfrak{X}(1)) = \pi_\Sigma(\mathfrak{X}(C))$.*

Proof. The non-obvious inclusion is $\pi_\Sigma(\mathfrak{X}(1)) \subseteq \pi_\Sigma(\mathfrak{X}(C))$, which follows from the Extension Axiom for fusion action systems and from the easy calculation that $C \leq N_{(\text{id}_1, \sigma)}$ for any $(\text{id}_1, \sigma) \in \mathfrak{X}(1)$. \square

Remark 3.5.2. We could therefore have defined transitivity of fusion actions in terms of the group $\pi_\Sigma(\mathfrak{X}(C)) \leq \Sigma_X$. If we want to concentrate on linking action systems, this alternate characterization has the advantage that C is always X -centric, and therefore we can define transitivity of a linking action system in terms of subgroups of S that are witnessed by $\mathcal{L}^\mathfrak{X}$.

We wish to introduce the notion of the “stabilizer” of a point $x \in X$ from the point of view of fusion and linking actions. Recall that we denote by \mathfrak{X}_x the stabilizer fusion action subsystem of \mathfrak{X} . Moreover, we shall denote by \mathcal{F}_x the underlying fusion system of \mathfrak{X}_x .

Definition 3.5.3. Given a linking action system $\mathcal{L}^{\mathfrak{X}}$ associated to \mathfrak{X} , the *stabilizer linking action system* of x , $\mathcal{L}_x^{\mathfrak{X}} = \text{Stab}_{\mathcal{L}^{\mathfrak{X}}}(x)$, is the category whose objects are those X -centric subgroups that are contained in S_x and whose morphisms are given by $\mathcal{L}_x^{\mathfrak{X}}(P, Q) = \{\mathfrak{g} \in \mathcal{L}^{\mathfrak{X}}(P, Q) \mid \ell_{\mathfrak{g}}(x) = x\}$.

Remark 3.5.4. We can think of \mathfrak{X}_x and $\mathcal{L}_x^{\mathfrak{X}}$ as the preimages under the natural maps $\text{Mor}(\mathfrak{X}) \rightarrow \Sigma_X$ and $\text{Mor}(\mathcal{L}^{\mathfrak{X}}) \rightarrow \Sigma_X$, respectively, of the subgroup $\Sigma_{X-\{x\}}$.

Again, we must restrict our attention to the stabilizers of *fully stabilized* points of X , namely those points whose stabilizers are of maximal order. The following is the linking action system analogue of Lemma 2.6.7.

Lemma 3.5.5. *The point $x \in X$ is fully stabilized if and only if $\widehat{S}_x|_C^C \in \text{Syl}_p(\mathcal{L}_x^{\mathfrak{X}}(C))$.*

Proof. The group $\mathcal{L}^{\mathfrak{X}}(C)$ naturally acts on X by the composition

$$\mathcal{L}^{\mathfrak{X}}(C) \xrightarrow{\pi_{C,C}} \mathfrak{X}(C) \xrightarrow{\pi_{\Sigma}} \Sigma_X.$$

Axiom (B) of linking action systems implies that

$$\widehat{S}_x|_C^C = \left(\widehat{S}|_C^C\right)_x,$$

and it follows easily from the definitions that

$$\widehat{S}_x|_C^C = \left(\widehat{S}|_C^C\right) \cap (\mathcal{L}_x^{\mathfrak{X}}(C)).$$

The core C is strongly closed in \mathcal{F} , so in particular it is fully normalized and $N_S(C) = S$. Proposition 3.4.9 then implies that $\widehat{S}|_C^C \in \text{Syl}_p(\mathcal{L}^{\mathfrak{X}}(C))$.

Thus we have again reduced the problem to the case of actual finite groups, as in the proof of Lemma 2.6.7, and the result follows. \square

With this interpretation of the stabilizer fusion, action, and linking systems, we find ourselves in the situation examined in [OV], and we recall the following result:

Proposition 3.5.6. *Let \mathcal{T} be an abstract transporter system associated to the fusion system \mathcal{F} on the p -group S . Fix a finite group Γ and a group homomorphism $\Phi : \pi_1(|\mathcal{T}|) \rightarrow \Gamma$, or equivalently, a map $\text{Mor}(\mathcal{T}) \rightarrow \Gamma$ that takes composition to multiplication and inclusions to the identity. For any subgroup $H \leq \Gamma$ let $S_H \leq S$ be the maximal subgroup whose elements (viewed as morphisms of \mathcal{T}) are sent to H , and assume that $S_1 \in \text{Ob}(\mathcal{T})$.*

Let $\mathcal{T}_H \subseteq \mathcal{T}$ be the subcategory whose objects are those of \mathcal{T} that are contained in H and whose morphisms are given by $\mathcal{T}_H(P, Q) = \{\mathfrak{g} \in \mathcal{T}(P, Q) \mid \Phi(\mathfrak{g}) \in H\}$.

Let $\mathcal{F}_H \subseteq \mathcal{F}$ be the fusion system on S_H generated by $\pi(\mathcal{T}_H)$, and let

$$\mathcal{T}_{S_H}^{\text{Ob}(\mathcal{T}_H)}(S_H) \xrightarrow{\delta_{H,H}} \mathcal{T}_H \xrightarrow{\pi_H} \mathcal{F}_H$$

be the restrictions of the structure maps for the transporter system \mathcal{T} . Then:

- (a) $\Phi(\text{Mor}(\mathcal{T})) = \Phi(\mathcal{T}(S_1))$.
- (b) \mathcal{T}_H is a transporter system associated to \mathcal{F}_H if and only if $\delta_{S_1, S_1}(S_H) \in \text{Syl}_p(\mathcal{T}_H(S_1))$.
- (c) If the condition of Item (b) is satisfied and all fully centralized $P \leq S$ have the property that $Z_{S_1}(P) \leq P$ implies $P \in \text{Ob}(\mathcal{T})$ then \mathcal{F}_H is a saturated fusion system.

- (d) If for all $P \in \text{Ob}(\mathcal{T})$ we have $P \cap S_1 \in \text{Ob}(\mathcal{T})$, then $|\mathcal{T}_H|$ has the homotopy type of the covering space of $|\mathcal{T}|$ with fundamental group $\Phi^{-1}(H)$.

Proof. [OV, Proposition 4.1]. \square

Proposition 3.5.7. *If $x \in X$ is fully stabilized, the stabilizer fusion system $\mathcal{F}_x \subseteq \mathcal{F}$ is saturated and $\mathcal{L}_x^{\mathfrak{X}}$ is a transporter system associated to it.*

Proof. From the natural map $\text{Mor}(\mathcal{L}^{\mathfrak{X}}) \rightarrow \Sigma_X$ and $H = \Sigma_{X-\{x\}} \leq \Sigma_X$, in the notation of Proposition 3.5.6 we have $C = S_1$, $S_x = S_H$, $\mathcal{L}_x^{\mathfrak{X}} = \mathcal{T}_H$, and $\mathcal{F}_x = \mathcal{F}_H$. That $\mathcal{L}_x^{\mathfrak{X}}$ is a transporter system associated to \mathcal{F}_x is then simply an application of Lemma 3.5.5 (b) to Proposition 3.5.6 (b).

That \mathcal{F}_x is saturated follows from Theorem 2.6.4, which implies that the stabilizer fusion action subsystem \mathfrak{X}_x is saturated. However, with the additional structure of a linking action system, the result is easier to prove, thanks to the work of Oliver and Ventura. Item (c) of Proposition 3.5.6 shows that if $P \leq S$ is fully centralized and $Z_C(P) \leq P$, then P is X -centric, i.e., that $Z_C(Q) = Z_S(Q; X) = Z(Q; X)$ for all Q \mathcal{F} -conjugate to P . If $\varphi \in \text{Iso}_{\mathcal{F}}(P, Q)$, then $\varphi(Z(P; X)) = Z(Q; X) \leq Z_S(Q; X)$. Because P is fully centralized, it is also fully X -centralized, so the assumption that $Z(P; X) = Z_S(P; X)$ and comparison of orders implies that $Z(Q; X) = Z_S(Q; X)$. Thus the conditions of Item (c) are satisfied and \mathcal{F}_x is saturated. \square

We now find ourselves in the following situation: Let $\mathcal{L}^{\mathfrak{X}}$ be a linking action system associated to the transitive saturated fusion action system \mathfrak{X} , and x a fully stabilized point of X . We would like to understand the topological information of the stabilizer linking action system $\mathcal{L}_x^{\mathfrak{X}}$ as it relates to that of $\mathcal{L}^{\mathfrak{X}}$, and indeed Proposition 3.5.6 gives us some relevant information in terms of subgroups of $\pi_1(|\mathcal{L}^{\mathfrak{X}}|)$. We can also calculate the homotopy type of $|\mathcal{L}_x^{\mathfrak{X}}|$ directly, as follows:

Let $\iota : \mathcal{L}_x^{\mathfrak{X}} \rightarrow \mathcal{L}^{\mathfrak{X}}$ be the inclusion functor, and let $F : \mathcal{L}^{\mathfrak{X}} \rightarrow \mathcal{TOP}$ be the left homotopy Kan extension of the trivial functor $* : \mathcal{L}_x^{\mathfrak{X}} \rightarrow \mathcal{TOP}$ over ι . We already have the functor $\mathbb{X} : \mathcal{L}^{\mathfrak{X}} \rightarrow \mathcal{TOP}$ defined as part of the data of $\mathcal{L}^{\mathfrak{X}}$, and we would like to relate these.

Proposition 3.5.8. *F is equivalent to \mathbb{X} as functors $\mathcal{L}^{\mathfrak{X}} \rightarrow \mathcal{TOP}$.*

Proof. Thomason's theorem [Tho] tells us that we have a homotopy equivalence

$$F(P) \simeq \text{hocolim}_{(\iota \downarrow P)} * = |(\iota \downarrow P)|$$

for all $P \in \text{Ob}(\mathcal{L}^{\mathfrak{X}})$. We first prove that every component of $|(\iota \downarrow P)|$ is contractible and that the components can be put in natural correspondence with the points of X .

The equivalence of F with \mathbb{X} is not natural: For every $y \in X$, Corollary 2.6.8 implies that $\mathcal{L}^{\mathfrak{X}}(S_y, S_x)$ is nonempty as x is fully stabilized. Let $\mathfrak{g}_y \in \mathcal{L}^{\mathfrak{X}}(S_y, S_x)$ be a choice of a morphism in this hom-set for each y , and assume $\mathfrak{g}_x = \text{id}_{S_x}$. Also let $\overline{\mathfrak{g}}_y = (\mathfrak{g}_y)_{\text{iso}}$ be the restricted isomorphism of \mathfrak{g}_y (cf. Corollary 3.2.13), \mathfrak{g}_y^P the restriction of \mathfrak{g}_y to $P_y \leq S_y$, and $\overline{\mathfrak{g}}_y^P = (\mathfrak{g}_y^P)_{\text{iso}}$ the restricted isomorphism of \mathfrak{g}_y^P .

For any $(Q, \alpha) \in (\iota \downarrow P)$, so that $Q \in \text{Ob}(\mathcal{L}_x^{\mathfrak{X}})$ and $\alpha \in \mathcal{L}^{\mathfrak{X}}(Q, P)$, we have $\ell_{\alpha}(x) = y$ for some $y \in X$. Since the pair $(c_{\alpha}, \ell_{\alpha})$ is intertwined and $Q \leq S_x$, we have $c_{\alpha}(Q) \leq P_y \leq S_y$. Thus we have the following commutative diagram in $\mathcal{L}^{\mathfrak{X}}$, where all morphisms labeled i are the obvious inclusions:

$$\begin{array}{ccccc}
 & & S_x & & \\
 & & \uparrow i & \nwarrow \mathfrak{g}_y & \\
 c_{\mathfrak{g}_y}(S_y) & \xleftarrow{\cong} & S_y & & \\
 \uparrow i & & \uparrow i & & \\
 c_{\mathfrak{g}_y}(P_y) & \xleftarrow{\cong} & P_y & \xrightarrow{i} & P \\
 \uparrow i & & \uparrow i & \nearrow i & \\
 Q & \xrightarrow{\cong} & c_{\alpha}(Q) & &
 \end{array}$$

α_{iso}

The top composition is α by definition. By the choices made above, we have $\ell_{\mathfrak{g}} \circ \ell_{\alpha}(x) = x$, and therefore there is a factorization

$$\alpha = \underbrace{\left(i_{P_y}^P \circ \overline{\mathfrak{g}_y^P}^{-1} \right)}_{\mathcal{L}^{\mathfrak{X}}(c_{\mathfrak{g}_y}(P_y), P)} \circ \underbrace{\left(\overline{\mathfrak{g}_y^P} \circ i_{c_{\alpha}(Q)}^{P_y} \circ \alpha_{\text{iso}} \right)}_{\mathcal{L}_x^{\mathfrak{X}}(Q, c_{\mathfrak{g}_y}(P_y))}.$$

This is the unique (as all morphisms of $\mathcal{L}^{\mathfrak{X}}$ are categorically mono and epi) factorization of $\alpha : Q \rightarrow P$ as a composite $Q \rightarrow c_{\mathfrak{g}_y}(P_y) \rightarrow P$, so what we have really proved is the following:

For any object $(Q, \alpha) \in (\iota \downarrow P)$ such that $\ell_{\alpha}(x) = y$, there is a unique morphism from (Q, α) to $\left(c_{\mathfrak{g}_y}(P_y), i_{P_y}^P \circ \overline{\mathfrak{g}_y^P}^{-1} \right)$ in $(\iota \downarrow P)$, namely $\overline{\mathfrak{g}_y^P} \circ i_{c_{\alpha}(Q)}^{P_y} \circ \alpha_{\text{iso}}$. In other words, we have found a terminal object in the component containing (Q, α) that depends only on our choice of the \mathfrak{g}_y and where α sends x . This shows that $F(P)$ is homotopically discrete, and as we have assumed that $\mathcal{L}^{\mathfrak{X}}$ is transitive, the components are naturally identified with the points of X .

All we have to do is see how F acts on morphisms, compared to the functor \mathbb{X} . Recall that \mathbb{X} sends the morphism $\mathfrak{h} \in \mathcal{L}^{\mathfrak{X}}(P, P')$ to the map of spaces $\ell_{\mathfrak{h}}$. On the other hand, $F(\mathfrak{h})$ is induced by the functor $(\iota \downarrow P) \rightarrow (\iota \downarrow P')$ that sends (Q, α) to $(Q, \mathfrak{h} \circ \alpha)$. If (Q, α) is in the component we have identified with y we have $\ell_{\alpha}(x) = y$, and then $(Q, \mathfrak{h} \circ \alpha)$ is in the component corresponding to $\ell_{\mathfrak{h}} \circ \ell_{\alpha}(x) = \ell_{\mathfrak{h}}(y)$. This is just to say that $F(\mathfrak{h})$ permutes the space homotopy equivalent X by the permutation $\ell_{\mathfrak{h}}$, so the result is proved. \square

Theorem 3.5.9. *In the above situation,*

$$\text{hocolim}_{\mathcal{L}^{\mathfrak{X}}} \mathbb{X} \simeq |\mathcal{L}_x^{\mathfrak{X}}|.$$

Remark 3.5.10. The final piece of interpretation comes from thinking of the left hand side as the p -local finite group action-theoretic version of $EG \times_G G/H$, and the right hand side as $EH \times_H *$.

References

- [Alp] J. L. Alperin. Sylow intersections and fusion. *J. Algebra*, 6:222–241, 1967.
- [Asc] Michael Aschbacher. Normal subsystems of fusion systems. *Proc. Lond. Math. Soc. (3)*, 97(1):239–271, 2008.
- [BCG⁺] Carles Broto, Natàlia Castellana, Jesper Grodal, Ran Levi, and Bob Oliver. Subgroup families controlling p -local finite groups. *Proc. London Math. Soc. (3)*, 91(2):325–354, 2005.
- [BK] A. K. Bousfield and D. M. Kan. *Homotopy limits, completions and localizations*. Lecture Notes in Mathematics, Vol. 304. Springer-Verlag, Berlin, 1972.
- [BLO1] Carles Broto, Ran Levi, and Bob Oliver. Homotopy equivalences of p -completed classifying spaces of finite groups. *Invent. Math.*, 151(3):611–664, 2003.
- [BLO2] Carles Broto, Ran Levi, and Bob Oliver. The homotopy theory of fusion systems. *J. Amer. Math. Soc.*, 16(4):779–856 (electronic), 2003.
- [Dwy] W. G. Dwyer. Homology decompositions for classifying spaces of finite groups. *Topology*, 36(4):783–804, 1997.
- [Gol] David M. Goldschmidt. A conjugation family for finite groups. *J. Algebra*, 16:138–142, 1970.
- [Lin] Markus Linckelmann. Introduction to fusion systems. In *Group representation theory*, pages 79–113. EPFL Press, Lausanne, 2007.
- [MP] John Martino and Stewart Priddy. Unstable homotopy classification of BG_p^{\wedge} . *Math. Proc. Cambridge Philos. Soc.*, 119(1):119–137, 1996.

- [Oli1] Bob Oliver. Equivalences of classifying spaces completed at odd primes. *Math. Proc. Cambridge Philos. Soc.*, 137(2):321–347, 2004.
- [Oli2] Bob Oliver. Equivalences of classifying spaces completed at the prime two. *Mem. Amer. Math. Soc.*, 180(848):vi+102, 2006.
- [OS] Silvia Onofrei and Radu Stancu. A characteristic subgroup for fusion systems. *J. Algebra*, 322(5):1705–1718, 2009.
- [OV] Bob Oliver and Joana Ventura. Extensions of linking systems with p -group kernel. *Math. Ann.*, 338(4):983–1043, 2007.
- [Pui1] Lluís Puig. Frobenius categories. *J. Algebra*, 303(1):309–357, 2006.
- [Pui2] Luis Puig. Structure locale dans les groupes finis. *Bull. Soc. Math. France Suppl. Mém.*, (47):132, 1976.
- [RS] K. Roberts and S. Shpectorov. On the definition of saturated fusion systems. *J. Group Theory*, 12(5):679–687, 2009.
- [Tho] R. W. Thomason. Homotopy colimits in the category of small categories. *Math. Proc. Cambridge Philos. Soc.*, 85(1):91–109, 1979.